



Whitepaper

Daide Rivola, Vasco Medici, Lorenzo Nespoli
Gianluca Corbellini, Davide Strepparava

January 16, 2018
v1.1

Contents

1	Vision	3
2	Hive Power Platform	4
2.1	Concept	4
2.2	Owner events	4
2.3	Mechanism of the decentralized Hive Power Energy Market	5
2.4	Beekeeper	6
2.5	State channels	8
2.6	Admin App	8
2.7	User App	9
3	Tokens	10
3.1	Hive Token (HVT)	10
3.2	Stable Tokens	10
4	Market Design	11
4.1	The Problem	11
4.2	The solution in a non flexible setting	11
4.3	The solution in a flexible setting	11
4.4	Data modification attacks	12
5	Use cases	15
5.1	Self-consumption Communities (SCC)	15
5.2	Micro-Grid	15
5.3	Distribution Grid	16
6	Business Model	17
6.1	Energy fees	17
6.2	Forecast service	17
6.3	Load management service	17
7	Roadmap	18
7.1	Hive Power Alpha (Q4 2017 - Q2 2018)	18
7.2	Hive Power Beta (Q2 2018 - Q1 2019)	19
7.3	Hive Power 1.0 (Q1 2019 - Q1 2020)	19
7.4	Hive Power 2.0 (Q1 2020 and beyond)	20
8	People	21
8.1	Team	21

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*

The power sector is facing a paradigm change, moving from a centralized structure with big power plants (hydro, coal, gas and nuclear) leading the energy market to a decentralized scenario adopting distributed energy resources (DER), such as solar and wind.

In this picture 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 a centralized planning and the increase of intermittent electricity production in the lower levels of the grid raise the stress on the electrical distribution grid and can lead to severe power quality problems, with which the distribution system operators (DSOs) need to cope. These problems are further emphasized by the increase of electricity consumption driven by the electrification of heat generation (heat pumps) and mobility (electric vehicles), which will tend to further increase power excursions in the distribution grid.

To overcome this problem significant investments in the grid infrastructure are expected. This could initiate a so called business “death spiral”, or “grid defection” [2, 3]. Prosumers reduce their regular power consumption from the central grid in favor of self-produced power. Besides, they can install batteries to further increase their energy independence. Reasonably, these customers will still depend on the central grid for emergency or peak use, so electric utilities will have to operate their costly infrastructure and power-generating capabilities even as revenues from consumption decline.

We believe this bleak scenario is not the future we want: a future with a widening gap between autarkic prosumers, almost disconnected from the rest of the grid, and simple consumers forced to pay for a more expensive and inefficient grid.

On the contrary, the future electrical grid will be characterized by an increase of exchanged energy between prosumers, consumers and electric utilities, optimizing the energy resource and the usage of the 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) could also participate to the energy market and thanks to their high flexibility they could quickly respond to dynamic price signals. This represents a big opportunity of cost reduction for the end user.

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 the participation, reaching at the same time a technical and financial optimum for the whole community. This is achieved by devising a (mathematically sound) market mechanism that incentivizes the participants to collaborate with each other, coordinating their production and consumption. In contrast to other energy exchange market schemes, the Hive Power platform takes also in account technical aspects, such as cables power rating and voltage limits in order to reach a multi-objective optimal solution. Hive Power is perfectly tailored to the current grid transition, enabling a safe and cost effective operation of the electrical grid by ensuring a fair and resilient energy market for all actors involved.

2 Hive Power Platform

2.1 Concept

A **Hive** provides a distributed energy market platform where every prosumer can exchange its energy consumption and production. The basic component of a Hive is the **Worker**. A Worker is a blockchain enabled electrical meter and it is generally linked to a prosumer, e.g. a single appliance, a single family household, a storage system, a solar power plant, an industrial facility. A Worker belongs to a single Hive. A Worker measures and certifies its own energy production and consumption, and joins the Hive energy market.

More precisely, the interface with the electrical grid is monitored by an **Ethereum Meter** equipped with a hardware oracle, a cryptographically attestable anti-tampering sensor. The certified energy data from the hardware oracle is gathered by a blockchain computer and tokenized by a smart contract. The Worker also includes a forecasting service that predicts the energy consumption/production and an energy bidding system interacting with the Hive through the worker smart contract. Each Hive is created by a **Beekeeper** and contains a **Queen**, which acts as an aggregator for the Hive. The main goal of the Queen is to help the Workers coordination, by collecting the production and consumption forecasts of the single Workers, aggregating them and communicating them back to the Hive in an anonymous form. The Queen also collects data about energy production/consumption related to the Workers and distributes the corresponding fiat to the Workers retaining a part for the Hive owner.

A Hive consists of a Queen, a population of Workers and an owner. Each Worker is logically linked to the Queen with a connector. Hives can exchange energy profiting of a hierarchical (mother-daughter) relationship between them. Each Hive possesses only one Queen, which in the upper hierarchical level will act as a Worker, representing the entire Hive as a single entity (Fig. 1).

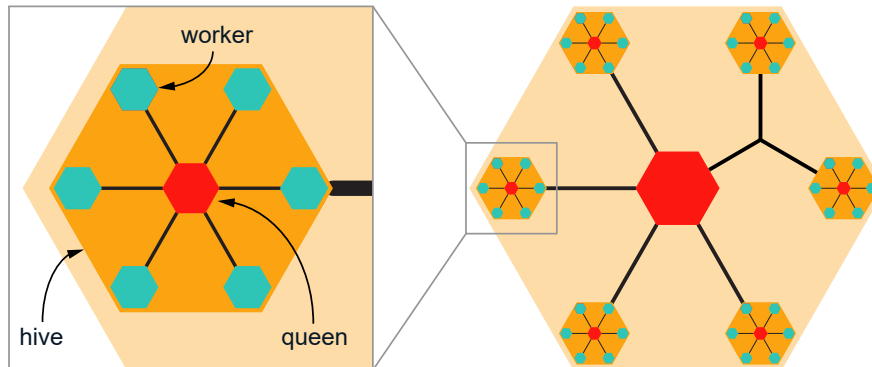


Figure 1: Hive components

2.2 Owner events

The interaction among a generic owner and the Beekeeper service is essential in order to correctly maintain a Hive. The main events occurring between an owner and the Beekeeper are explained in the following paragraphs.

2.2.1 Creating a Hive

The creation of a Hive is performed using Hive Token (HVT), an ERC20 token, as explained in Chapter 3. A generic actor (e.g. a DSO) using the HVT wallet requests a Hive creation for a group of meters to the Beekeeper service, staking and burning an amount of HVTs. The Beekeeper has to check if the HVTs amount is sufficient for the Hive creation with the number of requested Workers. It will also verify that the Ethereum Meters are certified for the usage in the Hive Power ecosystem. In case these requirements are fulfilled, the Beekeeper creates the Hive and assigns an identifier to it. The owner of the HVT wallet will be the **Hive Owner**.

2.2.2 Adding a Worker

A Worker (i.e. an Ethereum meter) can always be added to an existing Hive. First, the Hive Owner checks if the meter is listed in its own electric grid and has to stake new HVTs. Then - as in the creation event - the Beekeeper service checks if the new staked HVTs are sufficient and the meter is certified and the new Worker takes part in the energy market.

2.2.3 Dropping a Worker

A Worker can always be dropped by a Hive. Once dropped, the Worker leaves the energy market and a part of the staked HVTs goes back to the Hive Owner.

2.2.4 Splitting a Hive

A Hive Owner can request the Beekeeper to perform a split on one of its own Hives. The result of this operation is the partition of a Hive into multiple ones. After a successful split, the Hive Owner inherits the property of the new Hives.

2.2.5 Merging Hives

The Hives can also be merged. A Hive Owner has to request the Beekeeper to perform a merge on a set of its own Hives. After a successful merge, the Hive Owner inherits the property of the new Hive.

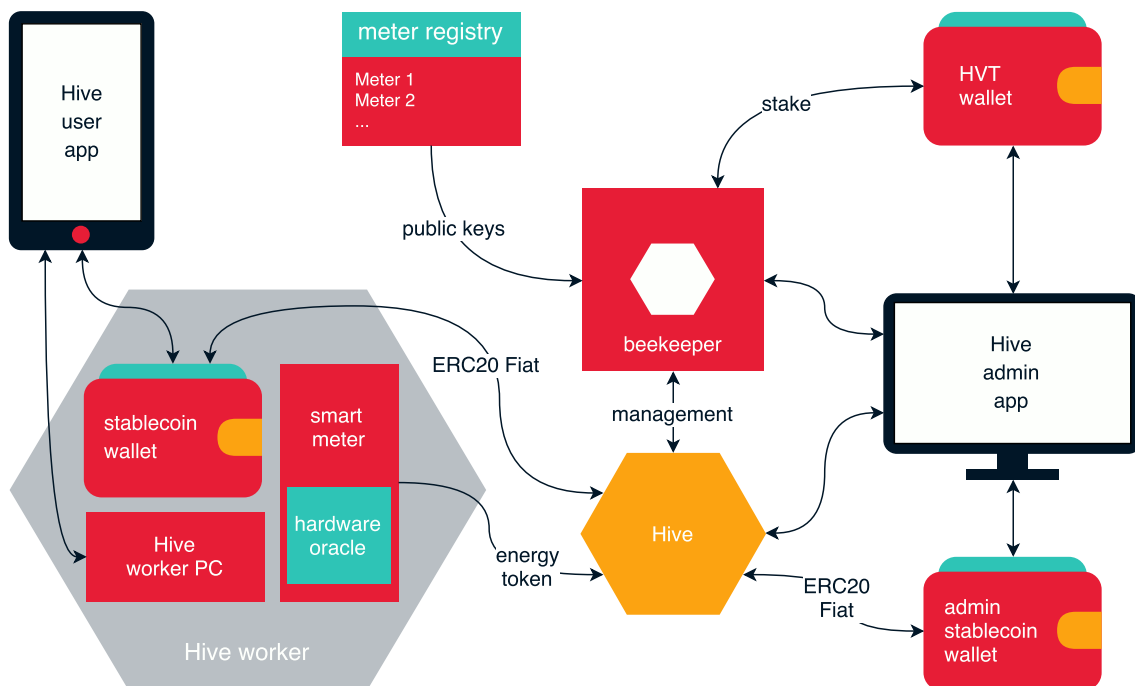


Figure 2: Hive Power platform

2.3 Mechanism of the decentralized Hive Power Energy Market

The **Energy Market** works with finite time steps (typically 15 minutes) and takes place in the following two distinct phases.

2.3.1 Coordination

In this phase, the Workers continuously send their forecasts to a Queen, who collects the data and communicates back to the Workers the aggregated profile and other useful market signals (e.g. prices). An agreement on the consumption and production profiles of the Workers in the following time step is reached in an iterative process.

2.3.2 Execution

At the end of each time step, the Workers communicate their certified consumption or production. Daily, the Queen collects this information and redistributes stablecoins to the Workers and to the Owner. The stablecoins had been previously staked by the Workers and Hive Owner.

In the case of a single Hive, the Hive Owner will receive enough stablecoins to be able to buy the electricity consumed by the Hive at the point of common coupling (PCC) with the rest of the grid, plus some additional fees, which cover the costs of the infrastructure and the power losses within the Hive. If the Hive is producing energy, the Owner would need to deposit enough stablecoins, to be able to pay the Workers for the energy they produce. In case the Hive is nested into a higher level Hive, the Queen will act as a Worker in the Hive of the higher level.

The Queen also verifies if the Workers actually diverge from their forecast values and assigns different reputations to them. This reputation is used to reward those Workers who better stick to their forecasts and punish those who are not.

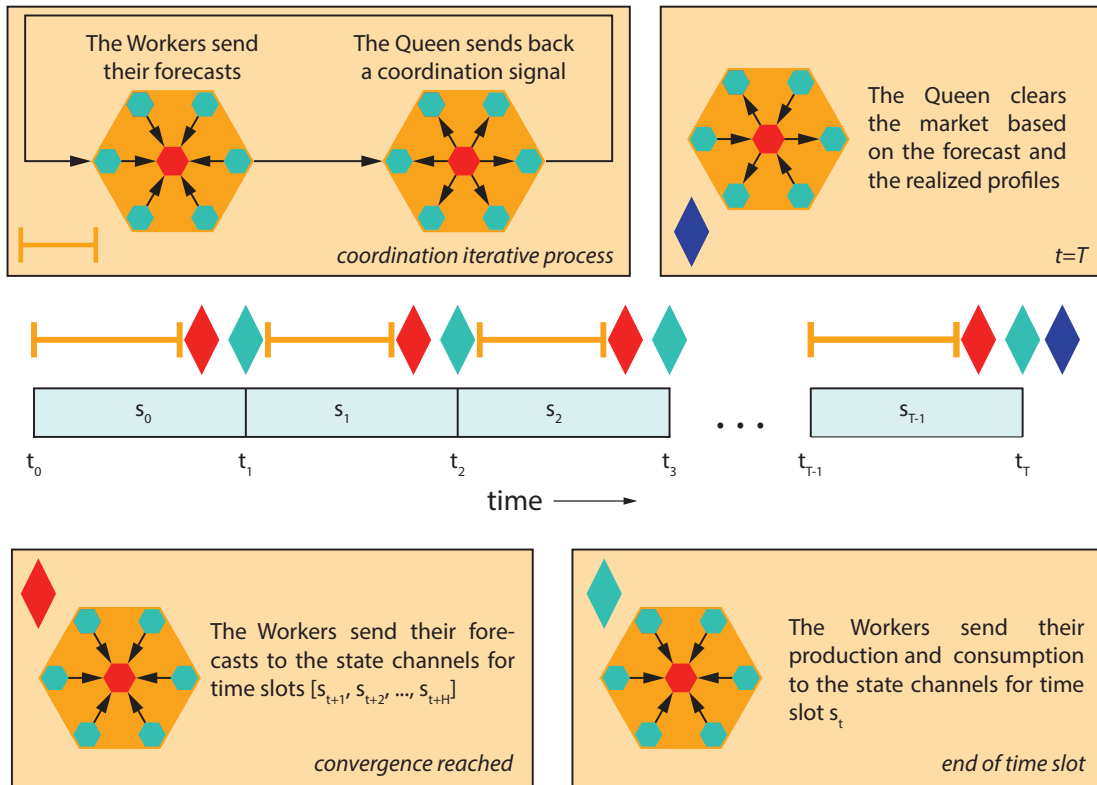


Figure 3: Hive Power market mechanism

2.4 Beekeeper

The purpose of the Beekeeper contract is to manage and track the Hives creation and deletion. Practically it maintains the DAO related to Hive Power platform. The main functionality of the contract is to provide three lists: the **Meter List** (ML), to track the certified meters available for the Hives; the **Stablecoin List** (SL), to track the stablecoins suitable for the payments and the **Hive List** (HL), to store data about the Hives themselves. The Beekeeper contract is able to add new entries in the lists and to edit the existing ones. On the other hand it does not delete anything, to maintain an exhaustive list of meters and Hives. Typically an HVT's owner can interact with Beekeeper adding/editing meters/Hives. For each created Hive, an amount of HVTs has to be staked by the owner in the contract. The Beekeeper contract is upgradeable in order to improve its security and to add new functionalities. An aim of Hive Power is to introduce in the future releases a governance system of the contract upgrades led by the HVT holders. Thus, HVT ownership acquires a further feature beyond the capability to create Hives, i.e. the capability to

address the development of Beekeeper and, as a consequence, to have a meaningful influence on the entire framework. In the next paragraphs the main features of HL and ML are explained in detail.

2.4.1 Meter list (*ML*)

To be able to create Hives, first at all the HVT owner has to insert its certified meters into *HL*. A meter always corresponds to a unique Worker in the Hive. For each inserted device the following mandatory features has to be initialized:

- Ethereum address of the meter
- Ethereum address of the meter owner
- Serial number of the meter
- Manufacturer
- Status (e.g. online/offline)
- Hive identifier

The meters list is upgradeable by the HVT owners (e.g. to change the status of a meter not available anymore). The exploitable functions for the interaction with *ML* are reported below:

```
addMeter(address meterAddress, address ownerAddress, string serialNumber, string
    manufacturer) returns (bool success)
changeMeterStatus(uint256 newStatus) returns (bool success)
getMeterInfo() returns (Meter meterData)
```

2.4.2 Stablecoin list (*SL*)

SL tracks the stablecoins suitable for the payments in the Hive Power platform. When a new Hive is created an available stablecoin is assigned to it. During the Hive life all the fiat transactions will be performed using the defined stablecoin. In case a Hive Owner wants to use a stabletoken not present in the list, it can request to the Beekeeper the insertion in *SL* of the desired token. In order to maintain an adequate security level the Beekeeper checks if the new token is eligible to be in *SL* and, in case, inserts it in the list. Periodically the Beekeeper controls *SL* and decides to “freeze” stablecoins if needed. A “frozen” stable token cannot be included in new Hives. The exploitable functions for the interaction with *SL* are reported below:

```
add(address stableCoin, string code) returns (bool success)
freeze(address stableCoin) returns (bool success)
getInfo() returns (StableCoin stableCoinData)
```

2.4.3 Hives list (*HL*)

As *ML* stores data about meters, thus *HL* has to track all the information about the Hives. When a new Hive is created, the following mandatory properties have to be initialized:

- Hive identifier
- Amount of HVTs for the staking
- Hive Owner address
- List of meters belonging to the Hive
- Status (e.g. online/offline)
- Starting timestamp
- Ending timestamp

- Stable ERC20 token used for the fiat transactions
- Status

Similarly to *ML* and *SL*, some features can be updated by the Hive Owner (e.g. a new certified meter has to be included in an already existing Hive, the Hive life is ended and it has to be “frozen”, etc.). The exploitable functions for the interaction with *ML* are reported below:

```

createHive(HVTToken amountToStake, address ownerAddress, Token fiatToken, Meter[] meters)
    returns (bool success)
freezeHive(uint idHive) returns (bool success)
getHiveInfo(uint idHive) returns (Hive HiveData)
addMeterToHive(uint idHive, address meterAddress) returns (bool success)
dropMeterFromHive(uint idHive, address meterAddress) returns (bool success)

```

2.5 State channels

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. To avoid this problem, for each meter only the final energy balance of a measurement period is stored on the blockchain. The energy measurements used by Hive Power to solve the market - with a typical resolution of 15 minutes - are managed using State Channels technology. This off-chain technology provides secure, fast and economic micro-transactions and can be automatically connected to Ethereum blockchain. The Hive Power alpha version is using an off-chain communication channel implementation for the energy tokenization. This is performed with an internal token used exclusively for exchanges of energy measurements within the Hive Power platform. Broadly, the operating mode of Hive state channels can be divided in the following three steps:

- At the beginning of a measurement period each Worker opens a state channel with its own Queen
- During the period, a data exchange between Worker and Queen is maintained by the related state channel
- At the end of the period, the state channel is closed by the Worker and the global balances are saved on the blockchain

The energy tokenization is not performed through a standard ERC20 token, but it is customized taking into account the framework specific features. During Hive Power development it is planned to migrate the current custom implementation toward scalable generic frameworks like Raiden Network (<https://raiden.network>).

2.6 Admin App

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

Beekeeper interface The owner is able to interact with the Beekeeper service for the general management of the Hive (e.g. creation, adding of a new meter, etc.).

Hive supervision The Hive Owner 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 Owner can supervise the status of its own stablecoin vault and perform a withdrawal.

2.7 User App

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

Cost savings Users are able to see the amount of local energy used from the Hive and the cost savings, compared to the regular utility tariffs.

Energy stats Users are able to see useful energy stats of his own Workers.

Credit fill-in Users can register a credit card to automatically top-up the Workers wallet with fiat currencies, using a third party exchange service. Using an ERC20 stable token for the fiat money will effectively protect the users against price volatility.

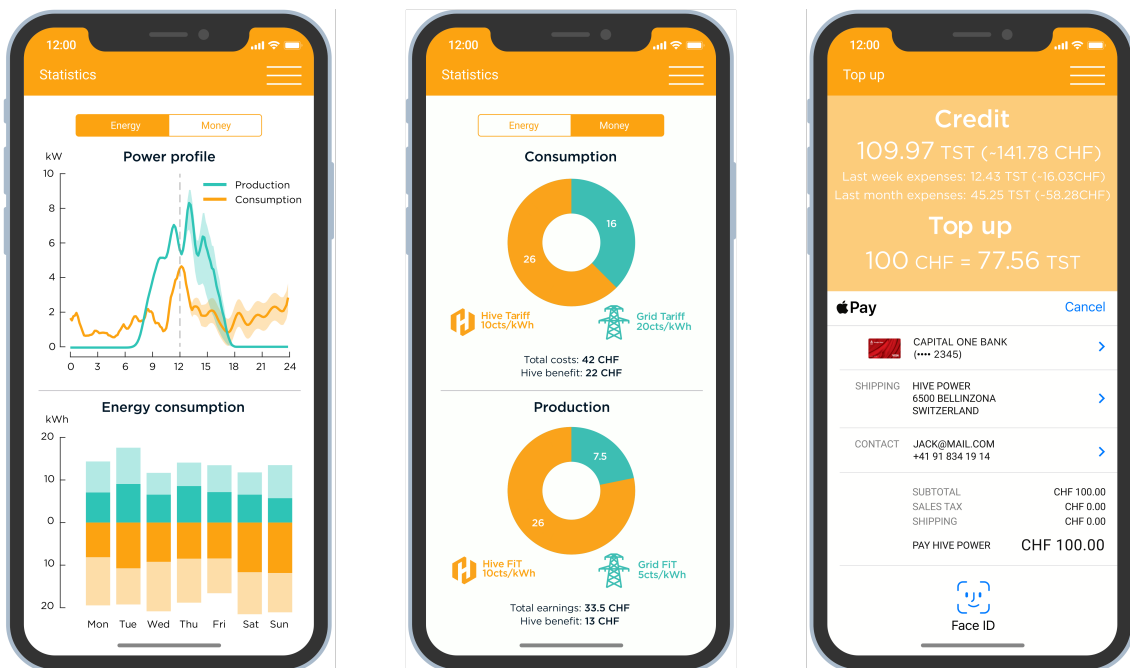


Figure 4: User App

3 Tokens

Two types of tokens are needed to operate the Hive Power platform. The HVT token is used to control and manage the operation of the Hives and to perform technical governance of the platform. The second type is used for payment operations, several third-party stable tokens can be used to fulfill this functionality.

3.1 Hive Token (HVT)

The Hive Token (HVT) is an Ethereum token managed by a smart contract. The purpose of HVT is to be used for creation and management of Hives and meters and participate to Hive Power platform governance. HVTs have a finite supply (100'000'000) and will be created only once, during the upcoming Token Generation Event (TGE). The Hive Token is a standard ERC20 token.

3.1.1 Hive management

A grid owner will need HVTs in order to create a new Hive. To do so the grid owner sends the required HVTs amount to the Beekeeper contract to perform these operations. The beekeeper contract perform a “burn and stake” on the received HVTs, 50% of the HVTs are taken permanently out of supply and the remaining 50% are staked inside the contract. The burn mechanism has the purpose to disincentive unnecessary Hive management operations and reach a stable operation of the Hive Power platform. After the creation of a new Hive, the Hive Owner can send additional HVT to the Beekeeper contract in order to attach meters to the newly created Hive. Also this operation follows the same “burn and stake” mechanism. When a Hive is destroyed or a meter is detached, the remaining 50% of the HVT tokens is sent back to the Hive Owner address. The required HVT amount for an Hive creation is variable in order to ensure a stable equivalent cost in fiat. The cost will be decided before the launch of Hive Power 1.0 release.

3.1.2 Governance

The Hive Owners, who are staking HVTs in the Beekeeper contract, will also have access to Hive Power technical governance. More specifically, all the upgrade proposals of Beekeeper contract will be subjected to the vote of Hive Owners. In order to exclude potential speculative actors, only Hive Owners will participate in the governance. Moreover the voting power will be weighted accordingly to staking age. Long term and active Hive Owners will have more voting power on the technical governance.

3.2 Stable Tokens

The Hive Power platform is designed as an open framework that can be used by any actor anywhere for multiple use cases. Due to its volatility Ether is not suitable to be directly used for the energy payments without integration of external price oracles, a complex solution. Therefore Hive Power platform payments will be operated via stable tokens, pegged to national currencies or any other kind of stable assets. The Hive Owner will decide during the Hive Creation event which stable token will be used in the related Hive. The interconnection of Hives using different stable tokens will be ensured by decentralized token exchange protocols such as <https://0xproject.com/> or <https://swap.tech/>.

Stable coin frameworks are in active development. DAI is the first platform we plan to integrate <https://makerdao.com/>. DAI is a ERC20 cryptocurrency pegged to USD (1 DAI = 1 USD) that automatically reacts to emergent market conditions in order to stabilize its value against the major world currencies. As the Ethereum ecosystem will mature additional stable coins will be added to the Hive Power platform.

4 Market Design

4.1 The Problem

To handle the indiscriminate injection of power from distributed energy resources (DER), DSOs currently use a simple indirect method, consisting in an asymmetric energy tariff, i.e. the prosumers buying price is higher than their selling one on the grid. This is a simple mechanism to increase individual prosumers self-consumption, but it is also rather inefficient in terms of their aggregated profile.

A better way to exploit prosumers' flexibility would be to explicitly formulate a common target for the aggregated power profile of the prosumers, and give them economic incentives to follow this target. Depending on the economic benefits, prosumers could choose how to use their flexibility.

For example, energy retailers, electricity suppliers or balance responsible parties, which bid for purchase of energy in the energy market, would benefit in a reduction of uncertainty in the prosumers consumption or production. Thus, they could forecast the aggregated power profile of a pool of prosumers and pay them to track this forecast. In this way, these energy actors obtain the same economic benefit of having a more accurate forecast.

Another revenue stream is represented by the possibility of creating self-consumption communities (SCC), which buys/sells energy as a single entity on the grid. Since in an 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.

$$\sum_i C_i \geq C_h \quad (1)$$

$$\sum_i f_c(E_i) \geq f_c\left(\sum_i E_i\right) \quad (2)$$

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

where \mathcal{C} and \mathcal{P} are the sets of consumers and producers at a given time, $E_i > 0$ representing the incoming (bought) energy for the i^{th} prosumer and $E_i < 0$ representing the outgoing (sold) energy. f_c is the cost function for Workers exchanging energy with the DSO, 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 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 a SCC. The economic surplus S guaranteed by Equation 2 cannot be used as a driver for the load coordination. Anyway, since $S \geq 0$, Workers will always have an economic benefit to participate to the market. The surplus S is simply redistributed applying a discount in the energy tariff in terms of [\$/kWh]. In this case there are obviously no benefits for the power quality in the Hive grid, nor for the grid on which the Hive interfaces, in terms of aggregated profile.

4.3 The solution in a flexible setting

If Workers are flexible, the surplus of money S can be used as a driver for the prosumers, pushing them to increase their self-consumption in order to maximize their welfare.

While the welfare maximization problem is conceptually simple to solve in a centralized way, in practice it leads to many technical and conceptual issues. The central controller needs to know all 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 EV charge. Gathering this information for all the Workers could be impractical and could raise concerns for the users privacy. Additionally, if grid constraints are taken into account, the central controller must retrieve grid topology and cables parameters. This coordination problem could become impractical to solve in real time when a large number of Workers is considered.

To solve this kind of issues, many distributed strategies for the 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. Distributed approaches can be classified in three main categories: decomposition techniques, networked optimization and non-cooperative games [5].

In the first phase, we will solve the coordination problem applying decomposition techniques. The Workers will locally solve an optimization problem in order to maximize the Hive welfare 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.

Workers solve their own problems in a declaration phase, using an iterative strategy. At each iteration, Workers communicate their solution to the Hive, which sends them back a coordination signal, allowing Workers to cooperate. The declaration phase ends when the iterative strategy has converged to the optimal consumption/production scheduling for the Workers. Since the Workers do not possess exact information about their future energy production and consumption, the declaration phase must be followed by an actuation one. The actuation phase starts just after the declaration and ends with the market clearance (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/produced, the surplus is redistributed and a new declaration phase starts.

Even if voltage and power constraints can be easily integrated in this approach, network topology self-discovery could be infeasible without high-frequency data acquisition or phasor measurement units. Thus, in this phase, 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, not being mapped to the physical network.

One of the core activities of Beta phase will be the extension of the proposed decomposition problem to a hierarchical setting. This will allow to create 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, optimization and communication strategies must be carefully designed in order to bound the required computational time.

The second core activity of Beta phase will be the automatic modeling of the electrical grid, starting from collected data. Having a model for the electrical grid will allow the Workers to aggregate their power profiles while respecting electrical grid constraints.

4.4 Data modification attacks

During the first phase, we assume that Workers solve the decomposed problem without cheating. This seems a reasonable assumption, since solving the decomposed problem maximizes the social welfare. Anyway, some malicious Workers could modify their local controllers, and assume a selfish behavior, trying to increase their own utility to the detriment of the other Workers. The selfish Workers will always try to fool the redistribution mechanism if possible, or to cheat revealing fake private information if this can increase their utilities. Game theory can be used to analyze this kind of opportunistic strategies, assuming Workers utility function is known.

In the following, we briefly introduce the game theory and decomposition approaches to the load coordination problem. We then analyze possible ways of fooling the described mechanisms and finally we describe a high-level description of the methods that will be implemented during Beta phase in order to mitigate these effects.

Dynamic Game Theory approach This approach can be considered as an indirect control, where we design the rules of an energy market for the prosumer community in order to induce a desired action (maximize self consumption). This energy market can be described as a non-cooperative n-person aggregative game [6] in a dynamic setting. Thus, we can use a game theoretic

approach in order to optimally design the market. Such design includes the choices of dynamic market prices as a function of user energy profiles and market clearance mechanism.

It is known that a non-cooperative n-person game leads to a unique Nash equilibrium if the utility functions of the users satisfy some mathematical conditions. If the game has a unique Nash equilibrium and its existence is common knowledge among the Workers, it is known that this equilibrium can be reached through a distributed algorithm [7]. This distributed algorithm, known also as gradient play [8], enables Workers to share their future energy plans with a central controller. The central controller aggregates the forecast of the Workers energy profiles and send it back to the Workers, which use this information to maximize their own utility. This allows cooperation between Workers, see for example the work in [9].

Decomposition and networked optimization approach The second approach is, on the contrary, a direct mechanism, where we explicitly formulate an objective function we want to minimize and we decompose it among the Workers. Decomposition techniques are commonly used to decrease overall computational time of a large optimization problem and to locally preserve private information; their properties are extensively studied by the scientific community. Networked optimization additionally considers communication constraints among the Workers: each worker can only communicate with a subset of neighbors. Even in this setting, if we know the grid topology, a coordination problem including grid constraints can be effectively solved [10].

4.4.1 Possible selfish strategies

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 in the game, Workers are allowed to cheat even in the presence of a Nash equilibrium. On the other hand, this aspect is very complex to model. Consider for example a market in which Workers pay a price proportional to their consumption and, if they consume more than the average, pay an additional fee. If Worker A declares a very big consumption for the following period, the other Workers could increase their consumption in the declarative phase, since they believe to be under-the-average consumers. If Worker A, after convergence, reveals that his needed consumption is in reality much less than the one he declared, he could now be a under-the-average consumer. The result is that he prevented the risk of paying the additional fee, to the detriment of the others.
2. If the Workers utility function does not generate a Nash equilibrium, it is likely that some Workers will try to solve a selfish version of the distributed problem, even if doing so they wouldn't maximize the welfare. This is true for all the approaches, not only for the game theoretic one. Consider for example the decomposition approach. If Workers utility function does not coincide with the decomposed problem, selfish Workers can choose to solve another problem. For example, a typical formulation of the decomposed problem leads to the following worker's optimization problem:

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

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}_ν 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}_ν 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

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

where $k \leq 1$, or simply lying on the true value of $\hat{\mathbf{P}}_\nu$. In this way he is prioritizing its own utility, forcing the other prosumers to take care of the grid constraints to the detriment of their economic benefits.

4.4.2 Proposed solutions

The first possible selfish strategy is a threat only when the number of Workers in the Hive is low or if the selfish worker is a big energy consumer/producer. Nevertheless, we plan to implement a reputation mechanism that prevents this kind of behavior. Briefly, 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 forecasting ability of the worker. Thus, with the introduction of this coefficient, we are at the same time promoting higher forecasting accuracy and discouraging selfish behavior.

In order to avoid the second selfish strategy, the worker optimization problem must reflect its own utility. Considering the decomposition approach, it is clear that the decomposed problem 4 does not reflect the economic utility of the worker, as previously explained, due to the first term.

Similar terms appear also in the game theoretic approach, in the presence of coupling constraints. One trivial way to force the Workers to solve problem 4 would be to make them pay for the first term. Anyway, this can incur in increasing energy tariffs for the Workers, who could opt-out from the energy market. A second option is to use different decomposition approaches to further decompose problem 4 into two subproblems, to be solved iteratively. In particular, the first subproblem can be represented by the minimization of the first and third terms of problem 4. Since these terms are known and common to all the agents, the minimization of this subproblem 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 $\phi(\mathbf{u}_i^\nu, \sum \mathbf{u}_{-i}^\nu)$, Workers must accept to cooperate with each other to fulfil grid constraints.

An alternative solution is based on detecting Workers selfish behavior, 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 market. From a probabilistic point of view, this can be seen as a conditional penalty on the worker's utility function.

5 Use cases

The Hive Power platform can be applied in several use cases, in order to optimize the behavior of Workers and the asset management of every node of the network. We distinguish three main cases, described in the following paragraphs.



Figure 5: 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 autarky**. Generators and consumers in a SCC must be on low voltage feeders on the same substation, they can internally optimize the synchronization of their production/consumption exploiting electric storage and demand side management. A simple case of a SCC is represented by a **condominium** where the solar energy, produced on the rooftop, is consumed by the tenants. In general, the members of the community, living in the same district, consume the produced solar electricity from their own roof. 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 inject the excess solar power to the national grid (e.g. during summer days) and get some remuneration. Batteries can be used to optimize the self-consumption rate of the SCC network. The clear economic advantage of a SCC comes from the large gap between the prices of energy bought (around 0.20 USD/kWh) and sold (around 0.05 USD/kWh) from the national grid.

5.2 Micro-Grid

As defined by US DOE, micro-grids are groups of interconnected 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 face a rapid deployment in the coming years as stated by the International Renewable Energy Agency in the IRENA Innovation Outlook 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 following different smart-pricing schemes:

- **Commercial and institutional (C&I) micro-grids**, aggregating existing on-site generation with multiple loads located in tight geography in which the owner is able to easily manage them. Particular cases are military bases and industrial micro-grids, where the power supply security and its reliability are very important. C&I micro-grid can connect and disconnect from the grid to enable it to operate in both grid-connected or island-mode. This case is

similar to a SCC, where the generation capacity is big enough to cover most of the local energy demand.

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

5.3 Distribution Grid

The third use case concerns the low voltage distribution grids, where a full district is connected to the national medium voltage grid, involving from few to hundreds of prosumers with different profiles:

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



Figure 6: Nodes of a Distribution Grid

The users can have generation systems, 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. All the users in the network take advantage from participating to the Hive, thanks to convenient tariffs (to buy and sell energy), so their assets (solar and batteries) become more profitable. In general, the adoption of favorable tariffs should significantly increase the penetration of DER (distributed energy resources) in districts with this market scheme, giving economic benefits to both producers and consumers.

6 Business Model

Hive Power is a technology provider of an innovative platform, fully open to existing and new energy actors. The goal is to keep 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 fees and additional services.

6.1 Energy fees

On each energy transaction performed inside Hive Power platforms, Hive Power retains a small fee, that is properly sized in order not to significantly affect the benefits of the stakeholders and collected during the settlement of the Hive Markets. This fee represents a compensation for the benefits that Hive Power brings to the involved 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 have a bonus tariff.
- Distribution System Operators (DSO) avoid the risk of user disaffection (death spiral) and, by using a prepaid billing system, they benefit from a reliable revenue stream.
- Hardware manufacturer can use Hive Power platform to add blockchain capabilities with reduced R&D costs and without paying licenses fees. Moreover, the deployment of SCC powered by Hive Power platforms can enhance the distribution of new hardware components.

6.2 Forecast service

In addition to basic functionalities, Hive Power will provide an accurate service for the forecast of energy consumption/production, used by Workers to optimize their own objective functions, piloting demand side management of batteries and thermal loads. A fee will be retained on the additional revenues generated to the worker owner.

6.3 Load management service

In the future, the Worker will gain the ability to control loads (such as heat pumps and boilers) and batteries to optimize its own performance (comfort and revenues). Load management features will also increase revenues from the Flexibility Market. A fee is required to activate the service.

7 Roadmap

The roadmap of Hive Power starts with the launch of a TGE, to finance the elaboration 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-community manager or a micro-grid manager).

First the focus will be on the development, testing and validation of the Hive Market design using an energy meter prototype, then we will finalize the whole 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 the already started discussion with potential partners interested to take advantage of Hive Power in their hardware and software solution.

7.1 Hive Power Alpha (Q4 2017 - Q2 2018)

The Alpha version will focus on a single self-consumption community, a rather simple use case even considering the existing regulation.

The building owner creates a Hive by staking HVTs and associates the Ethereum meters to the Hive. The apartment tenants buy fiat tokens with the Hive User App and use them to prepay the energy delivered to the apartment. By participating to the Hive, the tenants save money compared to apply regular tariffs. Building owners receive automatically the Hive payments and can use a part of that to pay the regular bill to the DSO. By using the Hive Power platform, the energy ledger is proven and tenants save administrative costs, avoiding bills and contract signings. The owner risk is neglected because tokens are prepaid by the tenants.

The Ethereum Meter proof-of-concept will run on the Hive Power Development Kit, a small open source computer connected to a triphase energy meter. As the goal is the testing of the market design, the payment channels will be implemented by 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 a first demo testbed, following called **Demo Hive** and shown in Figure 7, aiming to simulate a simple case of a Hive. More specifically, Demo Hive is constituted by two Workers (a producer and a consumer) and a Queen. Under the hardware point of view, the Workers are based on SmartPI boards, practically an acquisition board for the electrical measurements connected to a Raspberry Pi 3. Also the Queen device is a Raspberry Pi 3; primarily it has to manage the interactions between the Hive components and the external grid. Regarding the software, the main purpose of the demo testbed is to maximize the Hive autarky exploiting its tariffs, more convenient than the off-Hive ones. The basic behaviour is the following: at the end of a simulated day of 10 minutes, the produced/consumed energies of the Workers are calculated and tokenized using an ERC20 smart contract deployed on the Ropsten blockchain, the name of the related token is DHT (Demo Hive energy Token). In detail, the consumer worker owns a certain amount of DHTs and sends a part of them to the producers (i.e. the other worker, but also the external grid through the Queen if the consumption exceeds the production) in order to buy energy. Naturally, all the three devices are always connected and synchronized with Ropsten network to be able to transfer/receive the tokens.



Figure 7: Demo Hive testbed



Figure 8: Hive Power Alpha Development Kit

7.2 Hive Power Beta (Q2 2018 - Q1 2019)

At this stage the Hive Power platform will be fully operational regarding the basic functionalities. In this trial we will setup a complex Power Hive platform within a low voltage district in Switzerland. The Hive will consist of single household Workers and multi-family buildings, coupled to a low voltage grid. The Queen will be operated by the local DSO.

Ethereum Meter An Ethereum Meter prototype will be developed in collaboration with a meter manufacturer. The energy meter will tokenize exchanged energy with a certificate of origin. The energy token will not be exclusively tied to Hive Power. As tokenized energy will be a general use case that will be applied to other type of energy communities. The Ethereum Meter prototype will integrate a industrial system-on-chip. Lab tests are already on-going and we already verified that the chip is able to run an Ethereum light client, a state channel client and the worker software.



Figure 9: Current system-on-chip development kit

Load management The Worker will gain the ability to control loads (such as heat pumps and boilers) and batteries to optimize its own performance (comfort and revenues).

7.3 Hive Power 1.0 (Q1 2019 - Q1 2020)

At this stage the Hive Power platform with basic functionalities is validated and so available for the 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 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 Raiden payment network (<http://raiden.network/>). Deviation of agreed production and consumption are settled by a reactive control algorithm leveraging the Hive internal flexibility (storage, controllable load or generation) or through the bridging toward bigger and more flexible Hives. This 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 to the Hive's market.

Flexibility market In addition to the Energy Market inside and within neighboring Hives, it will also be possible to participate to a Flexibility Market. Hives could interconnect each other through virtual flexibility connectors and exchange spare flexibility. One will introduce the possibility to group Workers into virtual generating and consuming units, which will reach the required size to supply ancillary services to the grid, in particular primary and secondary frequency control, as well as congestion management. Load management feature will also increase revenues from the Flexibility Market.

Fraud detection Each element that is part of a Hive will stake a relevant amount of tokens. By comparing exchanged payment with real-time data, Hive components are able to detect fraudulent elements. Correction actions will include losing the staked amount, temporary or permanent ban from the Hive Power platform.

Hives flexibility Hives are not immutable, they are just virtual elements, therefore they can be reconfigured, merged and split.

7.4 Hive Power 2.0 (Q1 2020 and beyond)

After the release of version 1.0 many advanced features will be investigated:

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

Topology discovery When a new worker is activated it will discover autonomously the available Hives.

Self-healing A fault is automatically detected by the Hive and the involved component is flagged as defective. The Hive behavior adapts to overcome the current faults by activating spare flexibility.

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

8 People

8.1 Team

The founders of Hive Power have a wide academic and interdisciplinary background in the field of distribution grids/smart grids ranging from micro-technical engineering to applied mathematics and informatics. In the last seven years they gained important insights in developing an innovative algorithm based on a self-learning approach to the typical challenges of load management in smart grids. One of the largest utility in Switzerland has integrated this approach in a new product to be sold on the Swiss market by 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 a large experience in multinational corporations in the energy business, 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 modelling of photovoltaics plants, in the development of new business models for the optimization of smart grids. He was also lecturer for the 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, algorithms, 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 activity he worked for several years in building automation and access control industries. He has a notable experience in data science and database management. In SUPSI he is involved in research projects mainly related to the monitoring of solar plants and smart grids. In the last two years he matured an accomplished experience in blockchain technology, especially focused on 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 initiative 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] M. K. Jensen, “Aggregative games and best-reply potentials,” *Economic Theory*, vol. 43, no. 1, pp. 45–66, 2010.
- [7] J. B. Rosen, “Existence and Uniqueness of Equilibrium Points for Concave N-Person Games,” *Econometrica*, vol. 33, no. 3, pp. 520–534, 1965.
- [8] N. Li and J. R. Marden, “Designing games for distributed optimization,” *IEEE Journal on Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 230–242, 2013.
- [9] B. G. Kim, S. Ren, M. Van Der Schaar, and J. W. Lee, “Bidirectional energy trading for residential load scheduling and electric vehicles,” *Proceedings - IEEE INFOCOM*, vol. 31, no. 7, pp. 595–599, 2013.
- [10] M. Kraning, E. Chu, and S. Boyd, “Dynamic Network Energy Management via Proximal Message Passing,” *Foundations and Trends in Optimization*, vol. 1, no. 2, pp. 70–122, 2013.