Federated Learning on Trusted Data for Distributed PHM Data Analysis

Nikita Karandikar [1], Knut Erik Knutsen [2], Shuai Wang [3] and Grunde Løvoll [4]

[1, 2, 3, 4] *DNV AS, Veritasveien 1, 1322 Høvik, Norway*
nikita.karandikar@dnv.com
knut.erik.knutsen@dnv.com
shuai.wang@dnv.com
grunde.lovoll@dnv.com

ABSTRACT

Prognostics and health management (PHM) on systems such as vehicles and marine vessels are sometimes held back by complexities relating to data ownership and intellectual property rights. This is particularly true when multiple Original Equipment Manufacturers (OEMs) deliver components or sub-systems to a customer while having an interest in monitoring and maintenance of said component or sub-system. Further, the collection of PHM data from a fleet which may be non-uniform and spread across the globe with varying degrees of connectivity can be challenging from a bandwidth and cybersecurity point of view. Federated learning may address some of these challenges and open up new opportunities for how to approach PHM on a global and non-uniform fleet of components or systems. In this article we present FedChain, an approach for federated learning enabled by blockchain geared towards standardization for increased adoption. We discuss how a Docker based infrastructure for data collection, storage and analysis in combination with a methodology for tamperproofing PHM data can be a powerful substrate for bringing standardization, trust and transparency to federated learning implementations of PHM algorithms. We also demonstrate a basic blockchain enabled federated learning experiment and discuss the feasibility of applying FedChain from the perspectives of model performance, data privacy and security, tamper proofing and verifiability, and robustness.

1. INTRODUCTION

Prognostics and Health Management (PHM) of system components is increasingly employed in several industries, most notably in the aircraft industry (Xiongzi, Jinsong, Diyin, & Yingxun, 2011), and applications are developing rapidly. As an integrated technology, it utilizes sensor data and system information/knowledge to detect anomalies, diagnose occurring failures and predict the future development of failures, i.e. estimating the remaining useful life (RUL) of components/systems to support maintenance actions. In some cases such analysis is mandatory for safety reasons, such as Helicopter Usage and Monitoring Systems (HUMS) in the North Sea (CAA, 2014), while it is also considered an essential technology for improving maintenance efficiency and economics of operations. However, due to privacy and competitive considerations, sensor data are usually collected, stored and analyzed in a centralized manner depending on how the industry stakeholders are organized. In maritime, the prevailing notion is that ship owners also are the owners of data collected from systems on the vessels, while OEMs, who are often best positioned to analyze data and information from systems they have manufactured, usually have to seek permission to use such data, and data will usually not be shared between different OEMs. As such, the full body of data which can be of value to the industry as a whole, and in some cases outside a given industry, are fragmented between various ship owners as well as OEMs and often with severely limited availability to other stakeholders. Federated Learning, where various stakeholders may contribute and reap benefits without having to share all their data, which in some cases may be business critical, is a possible approach to improve models and advance PHM in maritime and in general. However, if the diverse data sources all have a single owner, such as in case of a fleet owned by a shipping company, then federated learning may not be necessary, a centralized model could be appropriate, and our approach would not apply.

Federated learning was introduced in 2017 by McMahan et al. (McMahan, Moore, Ramage, Hampson, & y Arcas, 2017) as a decentralized learning strategy to extract value from privacy sensitive or large sized data generated by mobile devices. Since then, there has been interest in applying this strategy of learning from sensitive data or data siloes in a variety of

different problem spaces such as health (Rieke et al., 2020), autonomous vehicles (Pokhrel & Choi, 2020) and industrial engineering (Hu, Gao, Liu, & Ma, 2018). Federated Learning enables utilization of data from multiple sources and data owners while preserving the privacy and intellectual property rights of each of these, to achieve powerful models so that each contributor may draw larger benefits than they would from only using their own data.

In this work, we propose FedChain, a federated learning approach enabled by blockchain, where the central server is partially or fully replaced by a smart contract and as such enables a transparent and distributed way of combining siloed data without sharing any raw data. We compare FedChain with centralised-based training approach and local-based training approach in terms of model performance followed by a discussion of the feasibility of FedChain from the perspectives of data privacy and security, tamper proofing and verifiability, and robustness.

2. DOCKER BASED INFRASTRUCTURE AND TAMPERPROOFING METHODOLOGY

The federated learning approach described in section 3 is designed for operating on a Docker (Docker, 2022) based infrastructure with the capability of sharing tamperproofed data both locally on the vessel and to shore. The modelling experiment described herein is done fully locally, but is designed to be easily deployable to Telenor Maritime's Docker based infrastructure called Unified Hosting Service (UHS) or to other similar architectures. The UHS infrastructure supports running of 3rd party docker containers, which are updated by the client.

A challenge arising when sharing data in the maritime industry is the lack of standard protocols and formats. Creating standardized exchange formats such as ISO19848 (ISO19848, 2018) and Maritime Context (Maritime Context, 2022) compliant data and establishing appropriate APIs would allow microservices bundled into Docker containers to easily exchange data and thus use multiple data sources for a particular application.These Docker containers do not have authority over the user's database, but only are able to access the data provided by the API. Docker is an open-source technology and a container file format for developing, shipping, and running applications. Docker allows users to separate applications in the infrastructure to form smaller particles (microservices), thereby increasing the speed of software delivery. The running environment of Docker images are standardized and can be deployed across different platforms. In order to manage the lifecycle complexity arising when many Docker containers are running, there is also a need for Docker Orchestration technology such as Kubernetes (Kubernetes, 2022), Docker Swarm(Swarm, 2022) as well as a GUI such as Portainer (Portainer, 2022).

Figure 1 shows the schematic of the tamperproofing methodology. Here, the developed microservices that offer various functionalities are shown in rectangles, while other commercial softwares are shown on the periphery. The arrows show the direction of data flow and the microservices are shown the the order of progression from left to right. The timeseries data is made available to the Data Collection microservice through the Telenor Maritime's API in batches. The microservice publishes this data to the MQTT (MQTT, 2022) message broker. The Hash microservice subscribes to the MQTT message broker and creates hashes for the timeseries data batches and publishes this hash to the blockchain, in our case VeChain (VeChain, 2022) and gets the transaction ID of the published hash. The hash and transaction ID are sent through Azure service bus (Azure, 2022) to the Hash MetaData microservice, which saves this data to Redis for future lookup. Finally, the Data Verification microservice creates a hash for the data to be verified and and looks up the transactionID for that hash in Redis (Redis, 2022). This transactionID is then found on VeChain and the hash there is compared to the generated hash. The rationale for using VeChain is discussed in section 5. The result of this comparison is returned along with the VeChain address and the time difference or lag between data generation as recorded in time series data and the time of publishing its hash to VeChain. A confidence score can be assigned to the data batch which is inversely proportional to the time lag. The older the data currently being hashed, the more time a malicious party has to tamper with it. The confidence score will be designed in order to reflect this uncertainty.

3. FEDCHAIN: FEDERATED LEARNING APPROACH ENABLED BY BLOCKCHAIN

Figure 2 shows the high level flow of FedChain, our proposed federated learning approach enabled by blockchain, which progresses through four stages. First, a global model is initialized. Second, a subset of nodes in the network are randomly selected to provide their local model parameters. Each of these nodes is given necessary access, to vote on the global model. Voting in this context means providing local model parameters that will be aggregated and thus influence the global model. As each node votes in a given round, voting access is revoked in order to restrict voting to those nodes who have been selected for each round and have not yet registered their vote. The global model is updated and is available for the next round.

Generally, in a federated architecture, a centralized server initializes the global model, chooses the contributor nodes for a round, accepts model updates, performs the aggregation and updates the global model. This server is also responsible for creating and administering the rounds and handling exchange of model parameters between the global model and each local model. However, such centralization can introduce concerns such as single point of failure and lack of transparency.

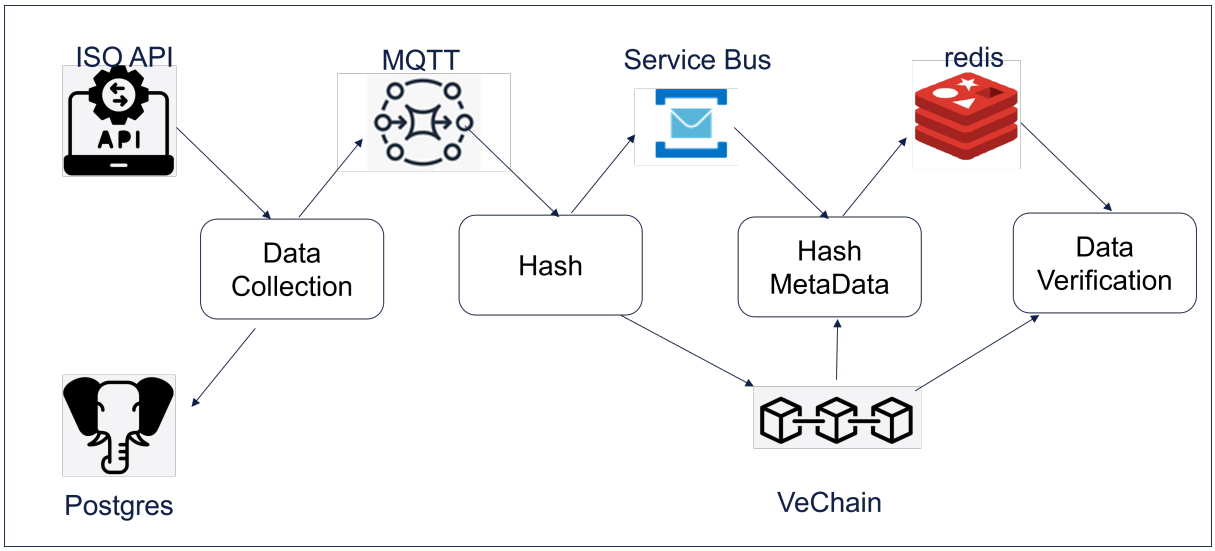

Figure 1. Conceptual illustration of tamperproofing methodology under development. By storing hashes of newly collected raw data as well as code or SW used for analysis/aggregation and any new data sets created, it is possible backtrace all changes to the data and to document the full data lineage. Any future consumers of raw or analyzed data can themselves verify that data remains tamperfree.

In order to address this, some works (Lo et al., 2021) have proposed using a blockchain for provenance tracing of global models with a centralized server which then performs aggregations and the exchange of local model parameters off chain. Such an architecture may reduce some of the transparency benefits introduced by using blockchain since the aggregation is still centralized. In effect, if the central server stops working the entire process halts.

To address the above-mentioned challenge, FedChain replaces a centralized server with a blockchain smart contract that provides a publicly available implementation for calculating aggregation. The smart contract state variables hold the global model parameters. Additionally, as identities on the blockchain are tied to their addresses, we use these identities to provide and revoke access to individual members of the federated network. As the blockchain smart contract is decentralized and replicated, there is no single point of failure or obvious attack surface for hackers. Further, as the implementation to calculate the global model is public, it allows for additional verification by federation members.

In a maritime federated learning scenario, we may have a large population of members. However, the potential contributors for each round may not be known in advance. Loss of connectivity, different update frequencies as well as possibility of dropouts are all considerations when selecting contributors. The consequence of this is that federated aggregations are expressed in a manner agnostic to the actual set of participants. Thus, the actual selection of participants, that may vary between training rounds is thus abstracted away from the computation and done by another component of the infrastructure. A distributed application (DApp) is an application that uses a blockchain as the backend, as opposed to a traditional application that uses a database. Such a DApp and smart contract based architecture can provide :

- a user friendly interface to access the blockchain functionality.
- trust in aggregation due to the publicly available smart contract code, with access to update the model restricted to privileged users.
- random selection of contributions in each round done off chain in order to reduce complexity and cost.

The address that deploys the smart contract becomes the federation chair. The roles, responsibilities and rights of the federation chair are different from the centralized server in a vanilla federation scenario. The federation chair can be the DApp if it is linked to a valid wallet. Algorithm 1 shows a high level view of the smart contract functionalities in FedChain. Each of these functions maps to a step of the process shown in figure 2, except the selection of the contributors for a given round, a task that would be taken over by the DApp. While the current aggregate global model is available at any time, periodically, an event, UpdateGlobalModel, is generated to inform the federation members that a cycle has ended and a new global model is created and should be updated.

4. EVALUATION

The goal of our study is to evaluate the proposed federated learning approach enabled by blockchain (FedChain) compared with 1) local model training (LMT), i.e., each client does not share any data or model updates and only train

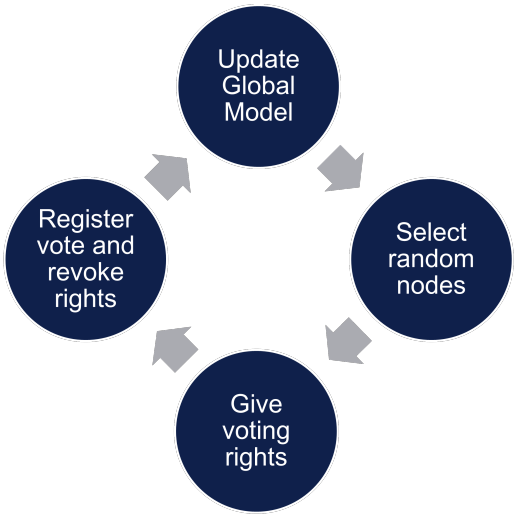

Figure 2. Overview of FedChain

Algorithm 1 FedChain

$coordinator \leftarrow msg.sender$;
procedure GIVEVOTINGRIGHTS($voter$)
Require: $msg.sender == coordinator$
$edges[voter] = true$;
end procedure
procedure REGISTERVOTE($local - parameters$)
Require: $edges[msg.sender] == true$
$global \leftarrow FedAvg(local - parameters)$;
$edges[voter] = false$;
end procedure
procedure UPDATEGLOBALMODEL
$emit$ updateGlobalModel;
end procedure

models with their own data locally; and 2) centralised-based model training (CMT), i.e., each client shares and updates all their training data to a centralised server, which trains model and distributes the updated model to each client for further usage (e.g., prediction). We first conduct a case study with the aim of comparing the model performance (measured by model score defined in Section 4.2) among FedChain, LMT and CMT. With such case study, we aim to address the following research questions:

RQ1: Does FedChain perform better than LMT?

RQ2: Does FedChain achieve equivalent performance compared with CMT?

Furthermore, we discuss the feasibility of FedChain compared with LMT and CMT from three perspectives in Section 4.4, i.e., data privacy and security, tamper proofing and verifiability, and robustness.

4.1. Data sets

In the maritime domain, batteries have been increasingly applied in both fully electric vessels and hybrid vessels in combination with combustion engines. As a result, understanding how battery capacity behaves over time (i.e., battery state of health estimation) is increasingly sought after (Vanem, Salucci, Bakdi, & Øystein Åsheim Alnes, 2021). This case study focuses on battery capacity estimation and employs laboratory battery cycle test data sets with the aim of evaluating the performance of FedChain, LMT and CMT. With these data sets, the objective is to train an optimal model (using FedChain, LMT and CMT) and estimate battery cell capacity with the model when a specific cycle number and temperature are given. More specifically, these data were collected from 27 battery cells and thus in total 27 sets of data were obtained. Each data set includes three columns, i.e., cycle number, temperature (i.e., ranging from 10°C to 40°C) and capacity measured with ampere hour (Ah). Figure 3 visualizes a data sample of the data set employed for cell No. 1, which we can observe that the battery cell capacity decreases when the number of cycle tests increases (e.g., the capacity drops to 59.908 Ah after the battery cell has been cycled with 480 times). Note that each battery cell is treated as an independent client when applying FedChain, LMT and CMT and thus we have 27 clients in total. Also, the data set for each cell is divided into two parts, i.e., 80% for the training data set and 20% for the test data set.

4.2. Model setting and evaluation mechanisms

To address the above-defined research questions, this study employs a linear regression model as the training model for FedChain, CMT and LMT and the formula for such model is shown as below. The goal is to obtain an optimal set of parameters (i.e., α_1, α_2 and β) for the model based on the

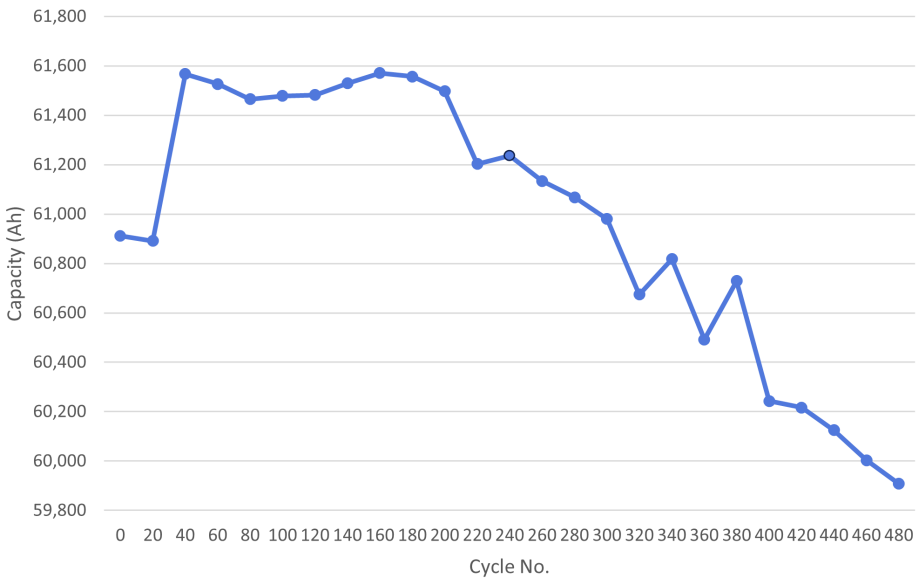

Figure 3. Data sample for cell No. 1 with the temperature 25°C

training data and estimate battery cell capacity with the model when a new cycle number and temperature are provided. It is worth mentioning that the various temperatures has been configured for cycle tests of different battery cells, e.g., 25 °C is set for the cycle test of cell 1 while 40°C is set for cell 5.

$$Capacity = \alpha_1 * CycleNum + \alpha_2 * Temperature + \beta \quad (1)$$

To evaluate the performance of FedChain, CMT and LMT, the metric model-score (MS) is employed (scikitLearn, 2019), i.e., MS returns the coefficient of determination of the prediction. More specifically, MS can be calculated using the following formula.

$$MS = 1 - \frac{u}{v} \quad (2)$$

where *u* is the residual sum of squares

$$\sum (capacity_{true} - capacity_{pred})^2 \quad (3)$$

while *v* is the total sum of squares

$$\sum (capacity_{true} - mean(capacity_{true}))^2 \quad (4)$$

Note that $capacity_{true}$ indicates actual capacity in the data set given a specific cycle number and temperature while $capacity_{pred}$ refers to the predicted/estimated capacity based on the trained model when inputting a cycle number and temperature. A higher value of MS demonstrates that the approach has a better performance to predict/estimate battery capacity when a cycle number and temperature is given and the best possible value for MS is 1, i.e., the approach can predict/estimate an accurate battery capacity for the entire training data set.

We have implemented FedChain, CMT and LMT based on two commonly-applied frameworks Flower (flower, 2022) and VeChain (VeChain, 2022) and a set of MS values (one value for each battery cell and in total 27 values) will be obtained for each approach. To compare the obtained results, the statistical t-test (Mcdonald, 2009) is employed to determine if there is a significant difference of the results (i.e., MS defined as above) between FedChain and CMT, and between FedChain and LMT. The significance level is set as 0.05, i.e., there is a significant difference between two approaches if the *p-value* calculated from t-test is 0.05 or lower.

4.3. Results

The figure 4 illustrates the achieved model-scores (MS) for FedChain, LMT and CMT, where we can observe that some of the local models performed worse than the other local models, e.g., linear regression model for battery cell no. 7 only managed to achieve 0.01 of MS with its own data locally. However, employing FedChain and CMT, the MS values for this cell were improved to 0.43 and 0.49, respectively, which indicates that the quality of the model was largely improved with the assistance of FedChain (sharing model updates) and CMT (sharing training data). From the figure, we also observed that linear regression model performed quite badly for at least five battery cells (i.e., cell no. 1, 7, 10, 17, 21) while these models have been largely improved when employing FedChain and CMT.

Furthermore, the average values of MS for LMT, FedChain

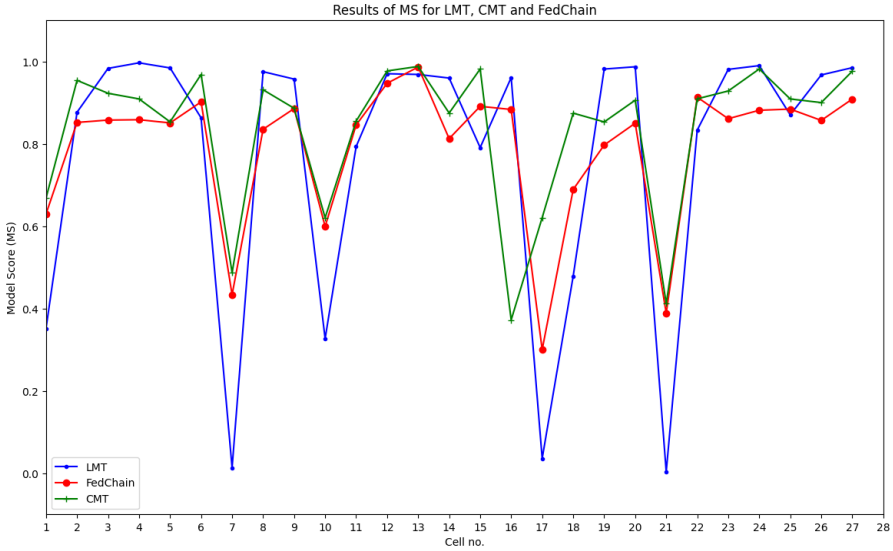

Figure 4. Results of MS for LMT, CMT and FedChain

and CMT are 0.774, 0.793 and 0.835, respectively, which indicates that CMT performed the best. We can also observe that FedChain managed to obtain higher average value of MS than LMT, showing that FedChain was able to improve the quality of trained models in general compared with LMT, especially for the cells performing badly (e.g., cell 1, 7, 10, 17, 21). The *p-values* of t-test between LMT and FedChain, and between FedChain and CMT are 0.79 and 0.39, respectively. The results showed that there was no significant difference of the MS results between LMT and FedChain, and between FedChain and CMT.

Based on these results, we can answer the two research questions as follows:

RQ1: As compared with LMT, FedChain can help to improve the model quality for certain number of cells while it can not significantly improve the model when considering all the battery cells;

RQ2: CMT achieved the best performance in terms of MS when compared with FedChain and LMT without significant difference, which indicates that FedChain can manage to obtain equivalent performance as CMT.

4.4. Discussions

In this section, we discuss the feasibility of FedChain compared with LMT and CMT in terms of the following three perspectives, i.e., data privacy and security, tamper proofing and Verifiability, and robustness.

Data privacy and security is of paramount importance in recent decades, where business/personal data should be well preserved and protected. However, such essential points have been neglected by traditional ML techniques (e.g., centralised-based ML approaches), where users' data requires being uploading to central servers for model training (Mothukuri et al., 2021). Federated learning (FL) (Aledhari, Razzak, Parizi, & Saeed, 2020) has been widely considered as a new paradigm of AI techniques where models are trained on client sides (rather than on central server side) so that users' sensitive data can be decentralised to avoid being shared among clients. Instead of directly sharing data, FL shares local ML model updates/parameters from each client with a global ML model that takes local models' updates as input for model tuning. When tuning is done, tuned updates will be distributed to each client for further model training. However, when local models communicate model updates with the global model, there is a chance that model updates can be tampered by 'unknown' organizations, which poses challenges for privacy and security for clients in terms of local models and even data. FedChain is proposed to address such challenges by integrating blockchain and FL, i.e., blockchain is used to ensure that model updates from each client or tuned updates from the global model can not be tampered anytime. Therefore, as compared with LMT (local training) and CMT (centralised-based training), FedChain offers a powerful means to preserve user privacy and enhance security by decentralising users' sensitive data and ensuring updates of local models tamper-proof.

Tamper proofing and Verifiability In FedChain, the model parameters are immutably stored in the blockchain, thus en-

suring non-repudiation and tamperproofing. Contributions are linked to the VeChain address, and aggregation code is open, so a contributor can verify that their contribution has changed the global model as expected. Similarly, due to provenance tracing offered by the blockchain, it is possible to verifiably trace the change in value. Spamming by participants is not a major concern here, as the participants are authenticated. Moreover, the right to contribute in a given round is explicitly assigned to the chosen participants and is revoked as the participant votes.

Robustness Learning collectively over diverse usage profiles will make models more generalizable. The global models , over several learning iterations get exposed to a significantly wider range of data that what is available to any one local model. As the models are trained collaboratively using a large amount of data, it also makes it more robust against single anomalous data points. However, centralization of data as in CMT is not only undesirable from a privacy point of view but also can use up available connectivity, a scarce resource for ship to shore communications. Moreover, as models are local, it makes real time prediction possible as the time lag resulting from transmitting raw data to a server and getting results back is eliminated. Local models continue to function even through loss of connectivity. A hybrid version of CMT could be considered where the centralized model is downloaded to function locally, but this configuration would still have the requirement of shipping all data over to the server, while not benefiting from locally generated real time data until such a time when the round trip of data and model parameters to and from the server is completed, relying heavily on the responsiveness of the server.

5. BARRIERS AND MITIGATION

Digitalization has seen an uptick in areas that traditionally relied on paper based documentation. The maritime sector has similarly shown increased interest in employing edge computing, data driven techniques and distributed ledger technologies for efficiency gains and cost savings in existing businesses and for creating new services (Sanchez-Gonzalez, Díaz-Gutiérrez, Leo, & Núñez-Rivas, 2019). Regulatory support for such digitalization has also been promising, as seen for instance in Singapore with the legal equivalence of electronic bills of lading (eBL) to traditional paper based versions (MPA, 2021a). Notably, TradeTrust (MPA, 2021b) is a digital utility comprising a set of globally accepted standards and framework to enable trusted interoperability of electronic documents, including eBL, through the integration of blockchain as an immutable ledger for record keeping, traceability and validation of such electronic documents. However, studies (Zhou, Soh, Loh, & Yuen, 2020) based on interviews with stakeholders have identified cost, complexity, lack of technical knowledge and unclear benefits as barriers for widespread adoption of such services. Federated learning relies on active collaboration between several participating entities, making adoption vital to the success of implementation. Thus our approach for federated learning enabled by blockchain will be discussed in the context of these barriers.

Consider a ship owner who would like to extract value from the data generated by on-board batteries. Such monitoring and analysis can can support planning of maintenance and replacement of batteries, performance modelling based on current State of Health or for safety purposes. However, the possibility of undetected tampering of the data, whether by dishonest or malicious actors or by accident, can severely undermine trust in the data and downstream services relying on it. Moreover, if this data is used for mandatory reporting or auditing to satisfy regulations, it must be possible for class societies to verify the provenance as well as veracity of the data. Blockchain can be used in order to bring verifiable trust to the data exchange value chain. However, there is technical complexity attendant in setting up such a solution for immutable tamper proof data for individual ships and their consumers. The containerized infrastructure described in section 2 would help develop a modular, extensible, standardized and replicable solution, thus alleviating complexity for the ship owner. Moreover, in order to implement a federated learning framework, such standardization would be vital for interoperability between multiple on board battery systems.

Cost was identified as a barrier, and potential for cost savings, a promoter for adoption of such a service. So, we now discuss the costs associated with the smart contract use. On public blockchain platforms, there are generally costs associated with transaction processing. This allows the platform to compensate the authority nodes for transaction processing and acts as a deterrent to parties who would otherwise inundate the platform with transactions. These costs can vary widely with factors external to the smart contract, such as with the platform in question as well fluctuations in price of tokens that must be spent to process transactions. Computation costs, storage costs as well as certain fixed costs depending on type of transaction may be incurred. High transaction costs and unpredictability associated with such costs can make it difficult to plan operating expenses and may be a deal breaker.

There are many public blockchains in existence but Bitcoin (Nakamoto, 2008), Ethereum (Buterin, 2014) and VeChain(VeChain, 2022) are among the most famous. On Bitcoin and Ethereum however, the token used to pay for transaction processing is also a publicly traded asset. Thus, the price of transaction processing in these platforms is highly volatile. VeChain, however has a two token system, wherein VET tokens are traded crypto assets and VTHO tokens are used to pay for transaction processing, thus stabilizing the transaction costs. VET holders generate VTHO, thus by holding VET, they are able to use the network for free and

sell excess VTHO for other users to purchase. Moreover, the immutability of transactions is one of the most attractive features of blockhain. However, the possibility of ledger forks calls into question the finality, thus immutability of posted transactions. Ledger forks are divergent ledgers that occur when participants of the network have a different view of finalized transactions. Both Ethereum and Bitcoin have had ledger forks in the past, in part due to their probabilistic consensus algorithms. Moreover, they rely on a highly computationally intensive consensus algorithm called Proof of Work, where nodes called miners race to solve a cryptographic puzzle to win the right to add the next block, expending tremendous amounts of electricity. For instance, Bitcoin mining consumes around 91 terawatt-hours of electricity annually, which represents 0.5 % of all electricity used globally (Jon Huang & Tabuchi, 2021).

VeChain, achieves consensus using a Proof of Authority based algorithm called SURFACE, an acronym for Secure, Use-case adaptive, and Relatively Fork-free Approach of Chain Extension. The key feature of SURFACE, as suggested by the name, is its focus on relatively fork free ledgers for enterprise scale networks. VeChain was also certified by the Centre Testing International Group Co. Ltd. (VeChain, 2021) as one of the greenest public blockchains in existance. Due to these reasons FedChain uses VeChain as the blockchain platform.

Transaction costs in VeChain are calculated in two stages. Equation 5 shows the formula for calculation of intrinsic gas cost $g_{intrinsic}$. This is the gas consumed by the transaction before any code runs. In simple terms, this is a constant transaction fee, plus a fee depending on transaction type, plus a fee for the size of input data for the transaction.

$$g_{intrinsic} = g_0 + g_{type} + g_{data} \tag{5}$$

$$g_{data} = 4 * n_z + 68 * n_{nz} \tag{6}$$

Here, g_0 is the constant transaction fee, currently set to 5,000 gas. g_{type} is 16,000 for a regular transaction, while it is 48,000 for contract creation. Finally, g_{data} is calculated as shown in equation 6. Here n_z is the number of bytes equal to zero within the data in the clause and n_{nz} the number of bytes not equal to zero. Beyond, intrinsic transaction costs, for a transaction to be run, a cost g_{vm} is incurred for the VeChain virtual machine that executes it. Moreover, as VeChain supports a multi clause transaction, equation 7 presents g_{total}, the total gas cost of a transaction execution, including multi clause transactions, where *i* is the number of clauses. Note, total cost also includes the intrinsic cost.

$$g_{total} = g_0 + \sum_{i=1}^{n} g^i_{type} + g^i_{data} + g^i_{vm} \tag{7}$$

Equation 8 translates the gas cost into gas price, where p^{base} is set to 1 VTHO per Kgas and ϕ is the value of field GasPriceCoef which is the bounded interval between 0-255. GasPriceCoef is used to adjust the priority of a transaction in the transaction pool.

$$p^{total} = p^{base} + p^{base} * \phi/255 \tag{8}$$

At the time of writing, 1 VTHO costs 0.001862 USD (Yahoo-Finance, 2022c), which is an order of magnitude lower than costs for Ethereum (Yahoo-Finance, 2022b) or Bitcoin (Yahoo-Finance, 2022a), thus making it a much more cost effective choice.

Thus, in this work we propose an integrated infrastructure for data management in order to ease implementation in real life scenarios, thus addressing the identified barriers relating to complexity, lack of technical knowledge and cost. Further, we use this trusted data for federated learning with value chain spanning tamper proofed architecture.

6. REMAINING CHALLENGES AND FUTURE WORK

From this work, several interesting research questions arise. Federated learning is organized into rounds where clients are chosen to contribute to a specific round. However, this can be complicated when clients have different update frequencies and granularity of data collection. Similarly, due to the nature of connectivity in marine environments, some clients may drop out temporarily. In these cases, the structuring and enforcement of rounds is challenging.

Model quality metrics can be difficult to establish without direct access to datasets. However, attribution techniques can be investigated, which can give an understanding of which contributions improve the global model. Based on this understanding, the federated averaging can be replaced by a weighted average, giving higher weights for better models or those trained on more instances. Incentivization and gamification techniques can be then incorporated to engage and encourage contributors. Such weighting can also be considered on the local model itself, to personalize the global model. Another possibility could be to cluster battery systems based on usage profiles, whether calculated or self-reported and create separate federations for each of these clusters. Finally, in order to prevent leakage of information about the training data set from the weights and parameters of the trained model, we can incorporate secure multi party computation and differential privacy. However, such techniques often have a high communication and computation cost, which may not be suitable for a maritime environment having to contend with limited data rates, low coverage or expensive setups.

REFERENCES

Aledhari, M., Razzak, R., Parizi, R. M., & Saeed, F. (2020).

Federated learning: A survey on enabling technologies, protocols, and applications. *IEEE Access*, *8*, 140699-140725. doi: 10.1109/ACCESS.2020.3013541

Azure. (2022). *Reliable cloud messaging as a service (maas) and simple hybrid integration.* Retrieved from https://azure.microsoft.com/en-us/services/service-bus/

Buterin, V. (2014). A next-generation smart contract and decentralized application platform. *white paper*, *3*(37), 2–1.

CAA. (2014). *Safety review of offshore public transport helicopter operations in support of the exploitation of oil and gas* (Tech. Rep. No. CAP 1145). Safety and Airspace Regulation Group, Civil Aviation Authority, Aviation House, Gatwick Airport South, West Sussex, RH6 0YR: Civil Aviation Authority.

Docker. (2022). *Developers love docker. businesses trust it.* Retrieved from https://www.docker.com/

flower. (2022). *Flower a friendly federated learning framework.* Retrieved from https://flower.dev/docs/index.html

Hu, B., Gao, Y., Liu, L., & Ma, H. (2018). Federated region-learning: An edge computing based framework for urban environment sensing. In *2018 ieee global communications conference (globecom)* (p. 1-7). doi: 10.1109/GLOCOM.2018.8647649

Jon Huang, C. O., & Tabuchi, H. (2021). *Bitcoin uses more electricity than many countries.* Retrieved from https://www.nytimes.com/interactive/2021/09/03/climate/bitcoin-carbon-footprint-electricity.html

Kubernetes. (2022). *Production-grade container orchestration.* Retrieved from https://kubernetes.io/

Lo, S. K., Liu, Y., Lu, Q., Wang, C., Xu, X., Paik, H.-Y., & Zhu, L. (2021). Blockchain-based trustworthy federated learning architecture. *arXiv preprint arXiv:2108.06912*.

Maritime Context. (2022). *Maritime context.* Retrieved from https://www.maritimecontext.org/about/

Mcdonald, J. (Ed.). (2009). *Handbook of biological statistics.*

McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics* (pp. 1273–1282).

Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. *Future Generation Computer Systems*, *115*, 619-640. doi: https://doi.org/10.1016/j.future.2020.10.007

MPA. (2021a). *Electronic bills of lading.* Retrieved from https://www.mpa.gov.sg/web/portal/home/maritime-companies/research-development/mint-fund-call-for-proposals/electronic-bills-of-lading

MPA. (2021b). *Tradetrust.* Retrieved from https://www.imda.gov.sg/programme-listing/international-trade-and-logistics/tradetrust

MQTT. (2022). *Mqtt: The standard for iot messaging.* Retrieved from https://mqtt.org/

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, 21260.

Pokhrel, S. R., & Choi, J. (2020). Federated learning with blockchain for autonomous vehicles: Analysis and design challenges. *IEEE Transactions on Communications*, *68*(8), 4734–4746.

Portainer. (2022). *Container management made easy.* Retrieved from https://www.portainer.io/

Redis. (2022). *A vibrant, open source database.* Retrieved from https://redis.io/

Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... others (2020). The future of digital health with federated learning. *NPJ digital medicine*, *3*(1), 1–7.

Sanchez-Gonzalez, P.-L., Díaz-Gutiérrez, D., Leo, T. J., & Núñez-Rivas, L. R. (2019). Toward digitalization of maritime transport? *Sensors*, *19*(4), 926.

scikitLearn. (2019). *Machine learning in python.* Retrieved from https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.LinearRegression.html

Standard data for shipboard machinery and equipment (Vol. 2000; Standard). (2018). Geneva, CH: International Organization for Standardization.

Swarm, D. (2022). *Swarm mode overview.* Retrieved from https://docs.docker.com/engine/swarm/

Vanem, E., Salucci, C. B., Bakdi, A., & Øystein Åsheim Alnes. (2021). Data-driven state of health modelling—a review of state of the art and reflections on applications for maritime battery systems. *Journal of Energy Storage*, *43*, 103158. doi: https://doi.org/10.1016/j.est.2021.103158

VeChain. (2021). *Vechainthor is one of the most eco-friendly public blockchains worldwide, cti verified.* Retrieved from https://www.vechain.org/vechainthor-is-one-of-the-most-eco-friendly-public-blockchains-worldwide-cti-verified/

VeChain. (2022). *Vechainthor.* Retrieved from https://www.vechain.org/

Xiongzi, C., Jinsong, Y., Diyin, T., & Yingxun, W. (2011). Remaining useful life prognostic estimation for aircraft subsystems or components: A review. In *Ieee 2011 10th international conference on electronic measure-*

ment & instruments (Vol. 2, pp. 94–98).

Yahoo-Finance. (2022a). *Bitcoin usd (btc-usd).* Retrieved from https://finance.yahoo.com/quote/BTC-USD/

Yahoo-Finance. (2022b). *Ethereum usd (eth-usd).* Retrieved from https://finance.yahoo.com/quote/ETH-USD/

Yahoo-Finance. (2022c). *Vethor token usd (vtho-usd).* Retrieved from https://finance.yahoo.com/quote/VTHO-USD/

Zhou, Y., Soh, Y. S., Loh, H. S., & Yuen, K. F. (2020). The key challenges and critical success factors of blockchain implementation: Policy implications for singapore's maritime industry. *Marine policy*, *122*, 104265.