electronics

MDPI

Article

Blockchain-Enabled IoT for Rural Healthcare: Hybrid-Channel Communication with Digital Twinning

Steve Kerrison [1], **Jusak Jusak** [1,*] **and Tao Huang** [2]

[1] School of Science and Technology, James Cook University, Singapore 387380, Singapore; steve.kerrison@jcu.edu.au

[2] College of Science and Engineering, James Cook University, Smithfield, Cairns, QLD 4878, Australia; tao.huang1@jcu.edu.au

* Correspondence: jusak.jusak@jcu.edu.au

Abstract: Internet of Things (IoT) and blockchains are enabling technologies for modern healthcare applications, offering the improved monitoring of patient health and higher data integrity guarantees. However, in rural settings, communication reliability can pose a challenge that constrains real-time data usage. Additionally, the limited computation and communication resources of IoT sensors also means that they may not participate directly in blockchain transactions, reducing trust. This paper proposes a solution to these challenges, enabling the use of blockchain-based IoT healthcare devices in low-bandwidth rural areas. This integrated system, named hybrid channel healthcare chain (HC^2), uses two communication channels: short-range communication for device authorisation and bulk data transfer, and long-range the radio for light-weight monitoring and event notifications. Both channels leverage the same cryptographic identity information, and through the use of a cloud-based digital twin, the IoT device is able to sign its own transactions, without disclosing the key to said twin. Patient data are encrypted end to end between the IoT device and data store, with the blockchain providing a reliable record of the data lifecycle. We contribute a model, analytic evaluation and proof of concept for the HC^2 system that demonstrates its suitability for the stated scenarios by reducing the number of long-range radio packets needed by 87× compared to a conventional approach.

Keywords: blockchain; digital twin; Internet of Things; healthcare; encryption; privacy; rural; LPWAN

check for updates

Citation: Kerrison, S.; Jusak, J.; Huang, T. Blockchain-Enabled IoT for Rural Healthcare: Hybrid-Channel Communication with Digital Twinning. *Electronics* **2023**, *12*, 2128. https://doi.org/10.3390/electronics12092128

Academic Editor: Djuradj Budimir

Received: 29 March 2023
Revised: 28 April 2023
Accepted: 4 May 2023
Published: 6 May 2023

1. Introduction

Internet of Things (IoT) technology has given rise to many new and innovative applications. In manufacturing, organisations from small to large scale use IoT to improve the monitoring of production processes, respond immediately when process deviation occurs, and to provide better services to their customers [1]. Implementation of IoT in the healthcare domain is a focus area for many researchers, academics, and industry as well. Healthcare IoT (HIoT) devices equipped with sensors, computation capability, and radio communications collect and process a patient's health related data, such as body temperature, electrocardiograph (ECG), oxygen saturation, blood pressure, and others to be transmitted to a cloud storage system in the other parts of the world through the internet. The term *Healthcare 4.0*, analogous to *Industry 4.0*, has been used widely to mark the development of smart and connected healthcare offering a chance to shift from traditional patient treatment to technology-based solutions that allow remote monitoring and medication [2].

Healthcare IoT is expected to be widely adopted but primarily benefits those in city regions who are most likely enjoying more extensive communications capabilities compared to those living in remote areas. The deployment of HIoT-supporting infrastructure in rural areas may face several obstacles. Geographical features of remote areas may be dominated by mountains, forest, savanna, hills, and rivers. In such areas, due to impediments to signals and low population density, there is less incentive for telecommunication providers to invest in installing significant infrastructure. Therefore, in most rural environments, low-communication quality, such as low bandwidth and intermittent connections, is frequently experienced by IoT devices, which can pose a challenge for real-time data usage.

Several technologies have been introduced in an attempt to address these adverse impacts, such as low power wireless area network (LPWAN) solutions [3,4]. LPWAN networks were introduced to accommodate the need for long-range and energy-efficient communications IoT devices. An example of such a technology is the long range (LoRa) standard that has growing adoption and industry support [5,6].

In addition to rural communication issues, HIoT faces security and privacy challenges in managing massive amounts of collected data. Cloud-based electronic healthcare records (EHR) emerged as a widely adopted solution [7]. They have several advantages, including on-demand service, broad network access, resource sharing, rapid elasticity, and guaranteed quality of service from service providers. With these features, the implementation of the EHR contributes to reduced data storage and maintenance costs, improved speed and processing accuracy, and allows data exchange among parties within a particular EHR system [8,9]. However, the centralised nature of the EHR system creates a setback from the user's point of view, in that users are more concerned about security and privacy due to the loss of control over clinical data in cloud storage.

Alongside the advancement of cloud and IoT, blockchain technology, the engine behind the cryptocurrency hype, has led to many other applications leveraging its features. For example, an article by Pennino et al. in [10] outlined the use of blockchain to support secure economic transactions underlying the decentralised payment system independent, the work by Wang in [11] investigated the utilisation of blockchain to secure energy delivery in electric vehicles, and some works by Farooq and Marbouh documented in [12,13] highlighted blockchain-based frameworks to assist healthcare management to monitor, diagnose, and treat patients remotely by stressing its applications in the most current COVID-19 pandemic situation. With the blockchain, certain aspects of applications become decentralised, in which control and decision are now shifted from centralised organisations to a distributed network. Each member node in a blockchain network retains a duplicate of the exact same information represented in the form of a distributed ledger. In this distributed network, consensus must be reached in order to add or change data, and the integrity of such operations is cryptographically verifiable. Attempts to tamper with information in the ledger is almost impossible.

In this work, we propose an integrated IoT and a private blockchain system applied to rural healthcare monitoring, called hybrid channel healthcare chain (HC^2). We chose a private blockchain scheme to facilitate a controllable environment, which is more appropriate for the healthcare use case than a public blockchain. The system operates two communication channels: short-range communication via personal area networks (PANs) for device authorisation and bulk data transfer, and long-range radio via LPWAN for light-weight data transmission and event notifications. Both channels leverage the same cryptographic identity information, and through a form of cloud-based digital twin, the IoT device is able to sign its own transactions via templates, without disclosing the key to said twin. Patient data are encrypted end to end between the IoT device and data store, with the blockchain providing integrity and authority only, thus protecting privacy.

The main contributions of the paper can be summarised as follows:

1. We define an architecture and data model for HIoT data that connects rural patients' data with healthcare providers with integrity provided by a blockchain.
2. We introduce a hybrid-channel communication model, allowing HIoT devices to use two communication methods to accommodate healthcare data transmission suitable for rural areas.
3. To overcome the transmission limitations on one of the two transmission channels, we incorporate a digital twin to handle data transactions from both of the communication channels and assist with blockchain transaction message reconstruction without sharing private encryption keys.

4. We demonstrate the benefits of our approach over the state of the art with a performance analysis based on the real-world constraints of LoRaWAN, a widely used LPWAN technology.

The rest of the paper is organised in the following order. We begin by discussing related works in Section 2, and proceed to provide a detailed description of our proposed model in Section 3. In Section 4, we present an implementation of the model using LoRaWAN and Hyperledger Fabric, with an evaluation and discussion of limitations of our integrated system in Section 5. Finally, we draw conclusions and discuss potential future work in Section 6.

2. Related Work

Our examination of related work begins with the challenges of rural healthcare monitoring, details the current technologies used for long-range communication, then looks at the uses of blockchain within healthcare, before summarising the combined challenges that we seek to address.

2.1. Rural Healthcare Monitoring

Providing appropriate communication infrastructure for electronic rural healthcare monitoring has been one of the most challenging issues from both the technological and economics points of view [14]. The geographical structure and population of these areas are the main reasons for this. Rural areas are often dominated by hilly terrain for large distances. Therefore, investing in the telecommunication infrastructure, such as 4th or 5th generation networks, in such areas has a low return on investment due to low population density and the complexity of installation for adequate coverage.

Alternatively, it has been suggested to exercise LPWAN technology, which lends itself to such settings due to low power transmission, while offering long-range communications among IoT devices. There are various standards bodies that are extensively working on developing LPWAN systems, such as the Institute of Electrical and Electronics Engineers (IEEE), the European Telecommunications Standards Institute (ETSI), the 3rd Generation Partnership Project (3GPP), the Internet Engineering Task Force (IETF), and the LoRa Alliance [15].

For example, a study by Dimitrievski in [3] showed the use of LoRa to carry healthcare data from rural areas combined with fog computing and the low Earth orbit (LEO) satellite connectivity to provide real-time data transmission. This work also proposed techniques for energy conservation utilising the external ultra-low-power timers that allow the device to be powered down, and showed its advantage to extend battery life in the order of tens of times. The fog system is a computation machine that is usually located between the cloud and the end devices to enable computing, communications, storage, and data management within the close vicinity of IoT devices. Therefore, in this IoT setting, the fog computation gives advantages to any delay sensitive devices to accumulate and process their retrieved data quickly (i.e., to achieve its real-time mode operation) rather than pushing through all data into the cloud system. Furthermore, the edge computing can be used to alleviate computing, storage, and bandwidth burdens of the system by allowing data processing within the edge devices when the resources of the IoT devices can be exploited to support that purpose [16].

Another study highlights a healthcare IoT architecture integrating blockchain and LoRa network to monitor patient health data securely [4]. To achieve real-time data transmission, the proposed model employs edge and fog devices to run the LoRa communication protocol whereby the edge devices with sensors attached on them collect data from healthcare data sources and subsequently send those relevant patient data to the upper fog layer using LoRa. To guarantee security, the data are stored in the interplanetary file system (IPFS) combined with blockchain technology. Finally, data monitoring and analytics for patients’ health status were performed through mobile or web applications.

The delivery of healthcare and the associated monitoring can be considered a complex system, with many changing variables that could change patient outcomes and affect decision making. The digital twin concept was first envisaged to aid the management of complex manufacturing systems, and the definition by NASA has become widely accepted [17]. Therein, a digital twin is considered a virtualisation of a physical system, maintained via the supply of data, for example, via IoT. With adequate data and modelling, scenarios can be simulated with a digital twin in order to predict outcomes for the physical system, allowing optimisations or corrections to be made. Unsurprisingly, this has been also been applied to healthcare settings [18]. In our work, we focus on the twinning aspects of HIoT sensors that allow the twin to facilitate blockchain-enabled activities that would not otherwise be possible over constrained network connections. As such, we assume that the wider benefits of digital twins (such as scenario simulation and physical/virtual linkages) can be realised elsewhere in the applications that make up the healthcare system as a whole. While we propose to use a twin to enable tighter integration between the HIoT device and the blockchain, a complementary (but not mutually exclusive) further example of their use can be in consensus-based decision making, such as that described for smart transportation, by Sahal et al. [19].

2.2. LPWAN and LoRaWAN

The term LPWAN, or low-power wide area network, refers to technologies that have the capability to reach long-range communications but at the same time maintain the minimum use of energy [6]. This communications model is particularly important to accommodate the need for various small devices which inherit features such as low computational power, low memory, and low battery capacity. However, contrasting these advantages of LPWAN, the nature of wireless signals dictates that most LPWANs have a low bit rate. Although there are many LPWAN architecture available on the market, LoRa has found its acceptance in both wider communities and broad industry support compared to other similar technologies in this scope, such as narrow band IoT (NB-IoT), LTE machine-type communication (LTE-M), and Sigfox [15].

Despite its long-range coverage and low-cost deployment, the most notable advantage of using LoRa is its reliance on a license-free operating frequency privilege operated on the industrial, scientific, and medical (ISM) frequency sub-band. The use of the chirp spread spectrum (CSS) modulation scheme on its bidirectional communications results in a signal transceiver with low noise levels yet high interference resilience. Utilising this modulation technique, the LoRa data rate varies from 250 bps to 50 kbps depending on the allocated spreading factor (SF) and channel bandwidth. For example, a lower spreading factor allows a higher data rate at the expense of a lower transmission range. The maximum payload length is 64–255 bytes, including its 13 bytes payload header, depending on the data rate chosen.

Alongside the growth of the LoRa adoption, LoRaWAN appeared as a protocol stack built on top of the LoRa physical layer. With its data link layer protocols support, this LoRaWAN shapes the LoRa network architecture into a typical gateway-nodes model that consists of a gateway that acts as a bridge between nodes, network servers and application servers over a backhaul interface [20]. In this structure, nodes can transmit messages to other LoRa devices or to a gateway. Hence, a gateway bears a task to gather data from all authorised sensor nodes (i.e., the end-devices) and pushes forward those data to the application server through the network servers.

The core of the LoRaWAN network resides in the network servers which maintain connectivity, routing, and security among devices. Therefore, gateways and network servers retain an important function in the LoRaWAN architecture to coordinate all nodes in its network, while at the same time synchronising data transmission to avoid collisions. This function was specifically defined in LoRaWAN as the medium access control (MAC) operation. Depending on how nodes should schedule their downlink traffic, users can alleviate the efficiency of LoRaWAN networks by properly selecting the class in which

LoRaWAN networks are deployed. The LoRaWAN allows operation in one of three different classes: A, B and C. In Class A (ALOHA) communications, an end device has the capacity to start transmitting data at any moment, whereas in Class B (Beacon), an end device can only open a receive window and transmit data between a periodic beacon signal duration according to the network-defined schedule. In Class C (Continue), an end device constantly listens to the downlink signal from the network unless the end-device is transmitting data.

Additionally, LoRaWAN enforces the duty cycle to limit the transmission of large amounts of data that may consume the whole bandwidth of a channel, which would cause congestion in the networks. The duty cycle defines how much of the total time a device is allowed to transmit data per hour on a particular sub-band. For example, a 1% duty cycle restricts the total amount of time a device spends transmitting data to 36 s per hour. Realistically, the amount of the duty cycle applied to a LoRaWAN is governed by regional regulatory authorities [21]. Furthermore, the things network (TTN), a service providing a public LoRaWAN network, applies a more rigid rule to lessen congestion by employing a fair access policy. This policy, applied to each end device, restricts the device's uplink airtime to 30 s per day (24 h) and downlink messages to only 10 in number per day [22].

2.3. Blockchain Systems for Healthcare

A considerable number of works have proposed IoT-based healthcare systems to provide a more timely and cost-efficient remote patient-care system [23,24]. Among other advantages, the IoT system might be identified as a substitute for the common in-hospital health monitoring with the remote one, where patients might stay at home or live in a rural area. While the traditional client–server and cloud computing paradigm offers significant improvement to the way patient data are stored, it also raises security and privacy concerns. For example, it suffers from the issues of single point of failure, data privacy, centralised data administration, and system vulnerability. The major threats to this cloud model may include spoofing identity, tampering with clinical data, and the data leaks [8].

Recently, the blockchain system has presented itself as a novel technology that could have a role in preserving healthcare data security and maintaining patient privacy. In a blockchain system, multiple data transactions, such as a patient's treatment and medical history, are grouped together in a structure called a block [25]. Each block is uniquely identified by its hash and timestamp and is chained to the previous block by incorporating the hash value of the previous block, thus creating a chain of blocks. The hash algorithm that is used acts as a one-way function, meaning it is computationally infeasible to produce a different block that would result in the same hash, effectively making the contents of the chain immutable. As such, validation of each block before they are chained in a blockchain network is paramount, as they typically cannot be removed or edited. Validation of transactions and blocks is performed by a consensus mechanism, whereby a shared ledger of blocks in the blockchain network can only be altered by the agreement or consensus of a majority of members [26].

Blockchain technology has a promising future in the healthcare domain, as it can solve some inherent issues facing modern health-management systems. It has advantages as a tamper-resistant distributed ledger for recording healthcare data and transactions, and its high availability and resiliency that will deter system failures and other cyber attacks [27–29]. However, the integration of blockchain into the IoT system in the healthcare rural area use cases may encounter several challenges to solve.

IoT devices may have difficulty to process and store even the smallest elements in the blockchain. Secondly, the geographical structure of rural areas and decreased availability of reliable transmission due to a sporadic communications infrastructure being in place are the other two notable problems faced by researchers to initiate such a secure healthcare monitoring system. As far as this study being carried out, we noticed there are only a few reported articles aiming to propose a solution in this domain. For example, the work by Munagala in [30] showed a blockchain-based traceable data sharing method to secure medical data transfer by incorporating software defined networking (SDN) technology to remove the clone nodes, and the work called Lorachaincare in [4] proposed a model of healthcare monitoring system which combines the blockchain, fog/edge computing, and the LoRa communications protocol. Besides focusing on its applications, there are also some blockchain-based frameworks proposed for managing secure healthcare systems, such as a framework for regulating mobile health apps and governing their safe use [13] and a framework for an asthma healthcare system that challenges its adoption during the COVID-19 pandemic [12]. All of these listed works use blockchain for medical data that have been collected in cloud storage, while the security of data transmission from IoT devices to the cloud is handled by encryption. However, the outlined works do not consider that transmitting medical data in rural settings is problematic, and several steps are required for user authentication in order to commit valid transactions in the blockchain system.

2.4. Use Case Definition

This paper seeks to further the state of the art by uniquely combining blockchain, LPWAN and HIoT technologies to deliver the possibility of improved healthcare services in rural areas. As such, we must address the following:

- HIoT data must be transmissible over an LPWAN technology that can be feasibly and cost-effectively deployed into rural settings.
- Integrity of data must be preserved through the use of blockchain, allowing lifecycle stages of the data (e.g., creation, storage, and granting of access) to be recorded.
- Patient confidentiality must be maintained, ensuring that persistent data such as those stored on the blockchain do not pose a privacy risk, nor are data transmitted over LPWAN a confidentiality or integrity risk if intercepted or manipulated.
- Mechanisms to provide the above security guarantees should be achieved alongside real-time transmission, avoiding the deferral of actions, such as the creation of transactions, wherever possible.

The following sections propose how to achieve these goals both architecturally and in implementation with the currently available technology, using the enhancements that we contribute.

3. Proposed Model

In this section, we describe our HC^2 model at a high level, and address healthcare entity participation, data flows, blockchain integration and security considerations. As an architecture, it does not dictate specific security, blockchain or communication technology selections, which we instead explore an example of in Section 4.

3.1. High-Level Architecture

Our HC^2 model uses Patel's framework for medical image sharing via blockchain [31] as a basis for its architecture. In the said work, image data are shared with the patient and physicians and forms the patient's health record (PHR), with access granted via transactions in the blockchain (as discussed in Section 2.4). First, we re-interpret this architecture to suit the HIoT use case, depicted in Figure 1.

The primary difference between this and the prior work is that the data provider is a healthcare IoT solution, rather than an imaging centre. The collection of data is not concentrated into a single location but rather streamed in real-time, or close to it, from a wide area, using many individual sensor devices. The HIoT data provider contributes sensor data to the patient's PHR and allows for any physician, authorised by the patient, to access them in order to provide them with healthcare services.

The sensor data are not stored on the blockchain, nor is any personally identifiable information regarding the patient. The access model for the blockchain is private, meaning that only authorised identities can view blockchain data and potentially transact on it. However, keeping personally identifiable information (PII) and sensor data off-chain provides additional protection of that data.

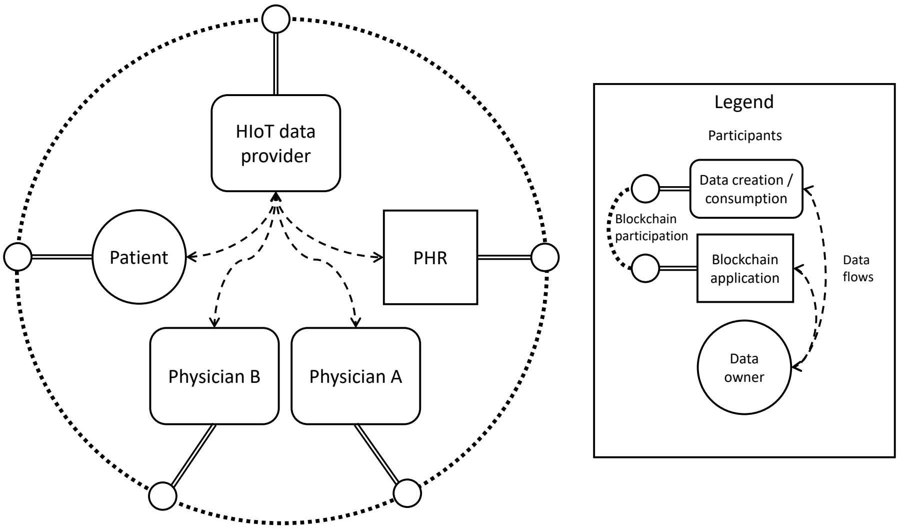

Figure 1. High-level architecture of HC^2.

3.2. HIoT Provider Entities and Data Flow

The HIoT provider component of the high-level architecture from Figure 1 comprises several entities that present unique challenges. We consider the following aspects:

- An HIoT sensor device which is paired with and attached to a patient for a duration of time. The device is expected to be portable and battery powered, for example, a health-monitoring watch or sensor pack.
- A "twin" of the HIoT sensor device used to represent the history and most-recent known state of the sensor device, regardless of connectivity status.
- Two communication methods between the device and its twin: one an LPWAN and one a PAN, where the LPWAN is low-bandwidth and possibly one-way, while the PAN is higher-bandwidth but intermittently available, for example, only when the patient visits a clinic.
- A data store for collected sensor and event data, obtained via the twin over either of the available communication methods.
- LPWAN connectivity is supported by base stations, uplinks to servers and subsequent internet connectivity to relay messages to the twin.
- PAN connectivity is achieved through short-range communication with an internet-connected bridging device, such as a phone over Bluetooth, or a physical docking station with USB or serial link.

The different types of participating components are represented with their own shapes. Potential data flows are represented by dashed lines, and the linkage to the blockchain, conceptually, is represented by the dotted circle around the diagram, to which the participants are all attached. Subsequent diagrams extend this concept further. For example, the participants responsible for maintaining the blockchain ledger and forming consensus are not represented at this stage.

The flow of data within this provision is visualised in Figure 2. The PAN is used for pairing, keying and detailed data transfers, whereas the LPWAN is used for small periodic data transmissions and events. For example, an HIoT device may monitor heart rate and ECG. After pairing, the device sends simple heart rate data over LPWAN every few minutes, along with an assessment of the patient's condition based on its own capabilities to analyse the heart rate and ECG data.

A healthcare provider may choose to act upon these data by calling the patient back to a clinic, or, under normal conditions, may await the next appointment. During the next visit, the detailed data logged by the HIoT device can be synchronised over PAN, via the twin, to the data store, and then immediately analysed for great insight.

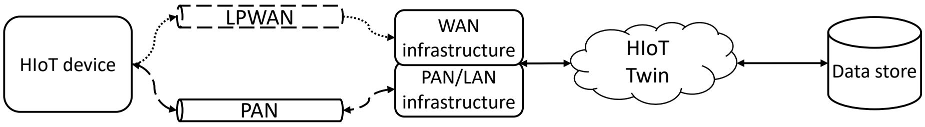

Figure 2. Hybrid-channel (LPWAN and PAN) architecture of HC^2.

3.3. Blockchain Integration

To integrate the HIoT provider into a decentralised blockchain, the following events must be recorded in the blockchain:

- Pairing between device and patient, whereby data that are generated by a device can be associated with the correct patient.
- Creation of data by the HIoT device to guarantee that a data record was produced by a legitimate, patient-paired source.
- Storage of data by an authorised data store to guarantee the retention of data that were generated by a device.
- Granting access to the data to additional entities to preserve a record of the management of permissions and, where necessary, encryption keys.

Pairing between the device and patient may be achieved through a transaction declaring the assignment of the device to the patient. The identifier for the device and patient must be sufficient to uniquely identify the relationship but does not need to be personally identifiable [32], and indeed this property may be necessitated by regulators now or in the future, who are advising on the best approaches to take [33].

While the HIoT provider may implement its own data store, this architecture does not preclude one or more external data stores being used, thus supporting a more decentralised approach to data handling. The data can be secured by a symmetric encryption key agreed between the sensor and store (discussed following subsection), and its creation, followed by its successful storage, recorded as transactions on the blockchain.

The events described above must be entered into the blockchain, and blockchain participants may refer to these in order to verify, authenticate, and progress to next steps in the process of providing healthcare. We focus mainly on the creation and storage of data in this paper (the middle two points), although all of these events can be considered blockchain transactions that must be recorded in a particular sequence in order for future actions to be allowed to proceed.

3.4. Security

The previous subsections alluded to several security considerations of the architecture, which we elaborate upon here. Firstly, sensor data are encrypted between the HIoT device and data store. To achieve this, a symmetric encryption method is used. If multiple data stores are used, then a key must be agreed between all of them and the device. To avoid overburdening the HIoT device, we assume that the data stores are responsible for coordinating key distribution among themselves.

This end-to-end encryption means that the device's twin cannot access the sensor data. It may store and forward the encrypted copy of the data, but will not possess the key needed to decrypt it. Data transferred over LPWAN or PAN are subject to this encryption, meaning the security of the WAN infrastructure or PAN link-layer poses no risk to the data's confidentiality.

The blockchain is largely responsible for protecting the integrity and availability of the data. Firstly, the creation of the data at the device is recorded as a transaction, verified by a signature that is cryptographically bound to the device's private key and associated identity. Similarly, the data store's acknowledgement of the receipt of the data has the same integrity assurances based on its own private key and identity information. Despite possessing the symmetric key used for data encryption, the data store cannot create data for itself, as it cannot sign a valid transaction representing the creation of data because its identity is not authorised to do so. Transformative processing of the data by other blockchain-enabled applications (for example, creating new data based upon analysis of the sensor data) remains possible and can be recorded as additional transactions, although the details of this are outside the scope of this work.

In terms of data availability, the loss of data over LPWAN can be established upon the synchronisation of data over PAN. At such a point in time, the device may verify that data it transmitted were correctly transacted, or the twin may observe the presence of records on the device that should have been received over LPWAN but were not. The cause may not be immediately knowable, but network outages, range issues or malicious interference can then be investigated. Finally, by agreeing on an expected data transmission interval, the twin may notify the HIoT provider system of missed data.

In summary, end-to-end encryption between device and data store provides confidentiality; blockchain transactions provide integrity and non-repudiation; the redundancy of communication channels (LPWAN + PAN) combined with the persistent presence/monitoring provided by the device's twin improves the detectability of availability issues; and the support for multiple external data stores improves the data's availability thereafter.

4. Implementation

To validate the architecture, we now discuss how it can be implemented in a realistic representation of our rural healthcare use case, under the constraints of contemporary HIoT devices, communication technologies, blockchain implementations and supporting software capabilities. First, we detail and justify our selections, then describe how the architecture can be realised within the technical constraints of the selections. Table 1 describes our selections, justifications for the choices, limitations/drawbacks and similar potential alternatives.

These technology selections pose challenges for how the components can fit the architecture of Figure 1 whilst achieving the requirements defined throughout in Section 3 in line with our use case. These are resolved in turn with the refinements detailed in this section, using proofs of concept where appropriate. Code for relevant proofs of concept, which are also used for data gathering used in Section 5.1, are collected into a group of repositories on GitLab [34].

Table 1. Technology selections made for HC^2 concept.

Component	Choice	Justification	Limitations	Alternatives
Sensor device	Micro-controller	Widely used for IoT-type devices. Relatively low cost. Capable of real-time sensor data acquisition.	Small amount of RAM and flash. Low processing power.	Smartphone or SBC with sensor attachments.
LPWAN	LoRaWAN	Multi-km range. Ability to create own infrastructure or use third party.	Limited or no downlinking. Very small uplink payloads and low duty cycles.	Narrow band IoT (NB-IoT), Weightless, Category M1 (Cat M1).
PAN	UART	Simplest communication method that can also be encapsulated within appropriate wireless protocols such as Bluetooth Serial Port Profile (SPP). Multi-kilobit to megabit transfer speeds are adequate for bulk data transfer.	Requires cable connection or dock to enable connection to twin.	Wi-Fi (LAN), ZigBee, Bluetooth low energy (BLE), serial peripheral interconnect (SPI), inter-integrated circuit (I2C).

Table 1. *Cont.*

Component	Choice	Justification	Limitations	Alternatives
Twin	Online deployment	Easier integration with LPWAN and connectivity to blockchain peers.	Link to device requires additional hardware with PAN + Internet capabilities. No offline capabilities.	Deployment onto LoRa gateways.
Blockchain	Hyperledger Fabric	Widely used for blockchain applications centring around business logic. Private access model. Certificate-based identities.	Transactions require round-trip communication with initiator.	Ethereum, Hyperledger Iroha.
Data store	MQTT historian	Commonly used protocol for IoT data simplifies collection of data. Multiple receivers can be implemented.	Blockchain application logic and data security must be additionally implemented and integrated.	Timescale, InfluxDB.

4.1. Blockchain Participation

We refer to the documentation for Hyperledger version 2.4, and, in particular, the “Key Concepts” topics, in describing the components relevant to this section [35]. Hyperledger Fabric uses public key infrastructure (PKI) to allow organisations to identify and enrol participants in the blockchain using certificates signed by certificate authorities (CAs). Fabric’s blockchain comprises several types of participant:

- **Committing peers** are responsible for maintaining the ledger state.
- **Endorsing peers** execute chaincode or smart contracts (the state-changing code executed with a transaction’s input arguments, described in [35]) to simulate a proposed transaction to determine if it would be valid.
- **Gateway peers** coordinate the dissemination of proposals to endorsing peers and the collection of endorsements on behalf of the proposer.
- **Orderer peers** construct blocks from endorsed transactions.
- **Clients** run applications that need to interact with Fabric peers to make transactions.
- **Admins** are able to perform privileged operations that change the configuration of Fabric and its peers, for example, by adding new organisations to a channel.

Identities for these participants are split into four groups: `client`, `peer`, `orderer` and `admin`. Most important to note is that not all participants maintain a copy of the blockchain or its current state. In the HC^2 model, this maps all participants of Figure 1 as *clients*, each possessing some form of application logic and an imperative to interact with Fabric to create, store and manage data.

Integrating the components of Fabric, alongside the other HC^2 components from Figure 2 into our architecture gives us a more detailed view, depicted in Figure 3. Here, we differentiate clients from peers. This creates a basis for visualising the sequence of transactions and flow of data in our rural HIoT use case.

The two biggest challenges from the limitations in Table 1 are the small packet size of LoRaWAN uplinks and the need for more than one round trip between client and Fabric peers to complete a transaction.

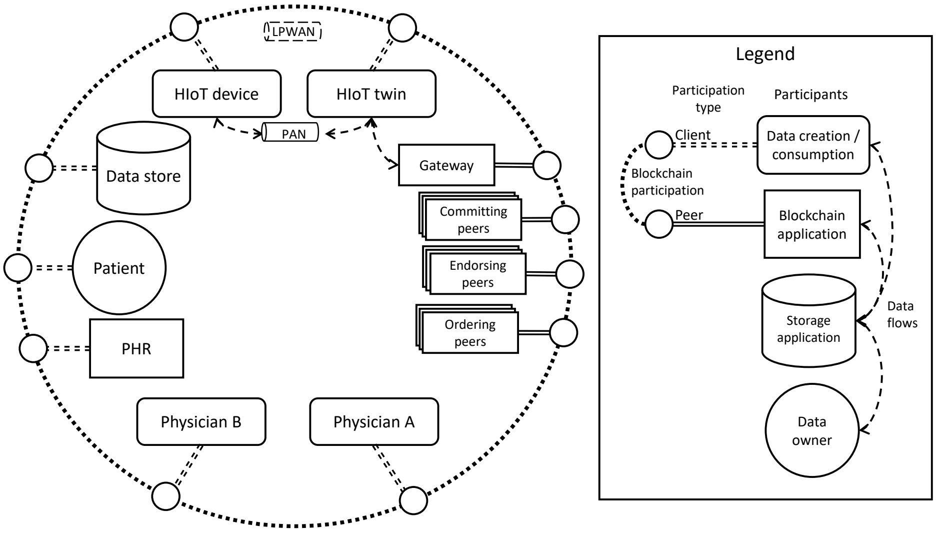

Figure 3. HC^2 architecture refined to accommodate selected technologies.

4.1.1. Payload Size

While various transmission profiles for LoRaWAN exist, the available payload size must accommodate the transmission of any data, in its encrypted form, along with a signature that might be usable in the blockchain. A 64-byte Elliptic Curve Digital Signature Algorithm (ECDSA) signature, as generated when using a P-256 (`secp256r1`) curve in Fabric, excludes most of the lower data rates from consideration. In the second Asia regulatory region (AS2 or AS923), which is of most interest to our research by virtue of locality, data rates providing 222 and 125 bytes of payload remain viable options [21].

Fabric uses `protobuf` to efficiently transfer messages between clients and peers. However, such messages still far exceed this payload limit, and when also faced with duty-cycle limitations as well, fragmentation is not practical.

We overcome this limitation by the pre-agreement of certain portions of Fabric messages, established between HIoT device and its twin over PAN (UART), prior to communication over LoRaWAN. Figure 4 shows agreements that take place between device, twin and data store. First, a template for Fabric messages is established between device and twin. Secondly, an encryption key is agreed between device and store, as discussed in Section 3.4, with the twin facilitating the transfer of the necessary key agreement messages. The fields and calculations that are relevant to the template agreement are detailed across Tables 2–4.

Table 2 lists the fields of a Fabric proposal that are agreed between the device and twin. This will be different for each device/twin pairing and each of their sessions but remain fixed between synchronisations.

The transmitted data are reduced down to that shown in Table 2, which is unique per transmission. Assuming messages may not be transmitted reliably and may not arrive in order necessitates the presence of a counter, C. These values are processed by the twin as indicated in Table 2 to complete the set of fields required to reconstruct the proposal message that was signed by the device.

Data transmission over LPWAN is reduced to three dynamic values: a counter C, encrypted data E, and a signature S. The latter two themselves are indexed by the counter value, and are all unique for each transmission. Within the 125-byte payload limit we selected for LoRaWAN, these values can be formed into a packet as shown in Table 5. The efficiency of this packet structure is discussed in Section 5.1.

From this packet, in combination with data prepared between the device and twin during PAN synchronisation, the twin is able to reconstruct the same message M that was signed by the device to produce its signature S. The twin can then submit this to the Fabric gateway on behalf of the device.

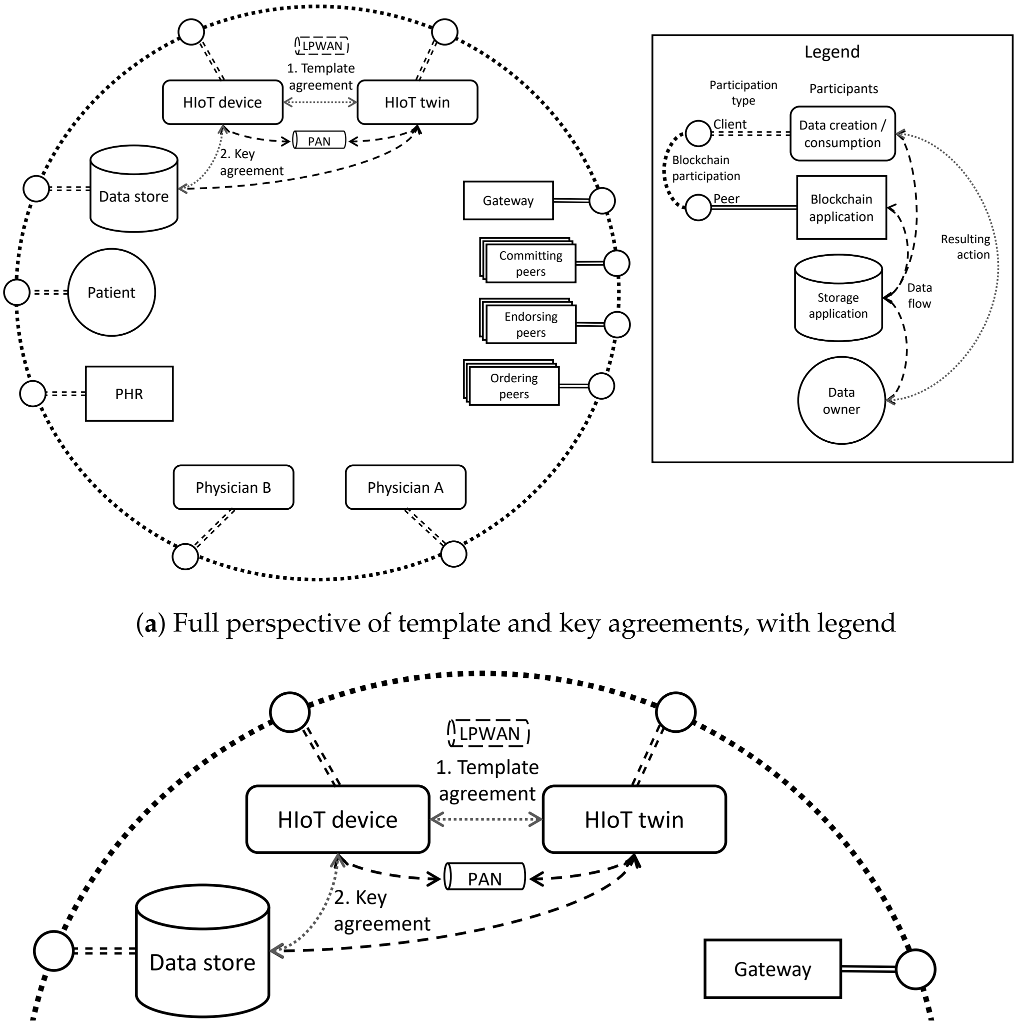

(**a**) Full perspective of template and key agreements, with legend

(**b**) Zoomed-in view of template and key agreements

Figure 4. Encryption key and proposal template agreement among device, twin and data store.

Table 2. Proposal fields agreed upon PAN synchronisation between device and twin.

Name	Symbol	Description
Header fields	H_x	Unchanging fields within the message header or headers of components within it.
Sync time	T_sync	Timestamp at synchronisation
Period	P	Time period stepped between transmissions
Identity	I_dev	Device's identity (certificate)
Seed	N_seed	Seed value used for per-transaction nonces
Args	$A_0 \ldots A_n$	Unchanging chaincode arguments

Table 3. Dynamic data sent by device over LPWAN.

Name	Symbol	Description
Counter	C	Number of messages since last synchronisation
Data	E_c	Encrypted sensor data
Signature	S_c	Signature of proposal as computed device-side

Table 4. Re-computed data by twin based on dynamic data and pre-agreed field values.

Name	Symbol	Computation	Description
Nonce	N_c	$N_{\text{seed}} + C$	Proposal nonce based on seed and counter.
Timestamp	T_c	$T_{\text{sync}} + PC$	Timestamp at which data was sent based on counter value
Data hash	$A_{\texttt{hash}}$	$\texttt{Hash}(E_c)$	A chaincode argument dependent on received encrypted data
Transaction ID	X_c	$\texttt{Hash}(I_{\texttt{dev}} \vert N_c)$	A unique ID for the proposed transaction based on nonce and creator.

Table 5. Payload format for HC^2 data over LoRaWAN.

Byte Position	0–1	2–65	66–124	***Total***
Length (bytes)	2	64	59	125
Purpose	Counter	Signature	Encrypted data	—

The precise construction of a transaction proposal is documented within Fabric's `protobuf` definitions [36]. However, referencing the values in Tables 2–4, we summarily describe the reconstructed proposal message in the partially abstracted form:

$$\begin{aligned} M = H_{\text{envelope}} | T_c | H_{\text{channel\_info}} | X_c | H_{\text{chaincode\_info}} | \\ | A_0 | \ldots | A_n | A_{\text{hash}} | H_{\text{signature\_info}} | S_c \end{aligned} \tag{1}$$

where the vertical bar symbol represents the concatenation of the values on either side of it, as an array of bytes. The exact ordering and encoding must respect that defined in the `protobuf` definitions [36].

4.1.2. Transaction Processing

Hyperledger Fabric ensures the integrity of transactions through endorsements. A client proposes a transaction, and the relevant chaincode is executed by several endorsing peers, and if valid, the peers sign and return endorsements to the client. The client can then combine these endorsements with the original proposal, signing them into a transaction which can be submitted, ordered and committed to the ledger, with the world state updated accordingly. In Fabric version 2.4 and above, the Fabric gateway can be used to distribute the client's proposal to necessary peers and collect the endorsement responses, prior to returning to the client for formation of the transaction submission.

This offloading is beneficial to the HIoT device, as it does not need to handle as much communication with Fabric. However, without refinement, transactions would only be proposed and endorsed but not committed, as the final endorsed transactions cannot be submitted to the orderer until the twin has an opportunity to return endorsements to the device, which we assume must happen over PAN. While data may be entered into the data store and the proposals/endorsements available in activity logs, this would delay the committing of any transactions to the blockchain.

To overcome this, we consider the chaincode for the data's early lifecycle in three parts:

1. The twin is responsible for submitting a transaction that *creates* the data.
2. The data store submits a transaction to register the *storage* of the data.
3. The device submits a transaction that *verifies* the data's origin.

It is counterintuitive to observe that the twin is responsible for the first transaction while the device is responsible for the last. Figure 5 shows the sequence of communications leading to transactions that achieve the desired outcome.

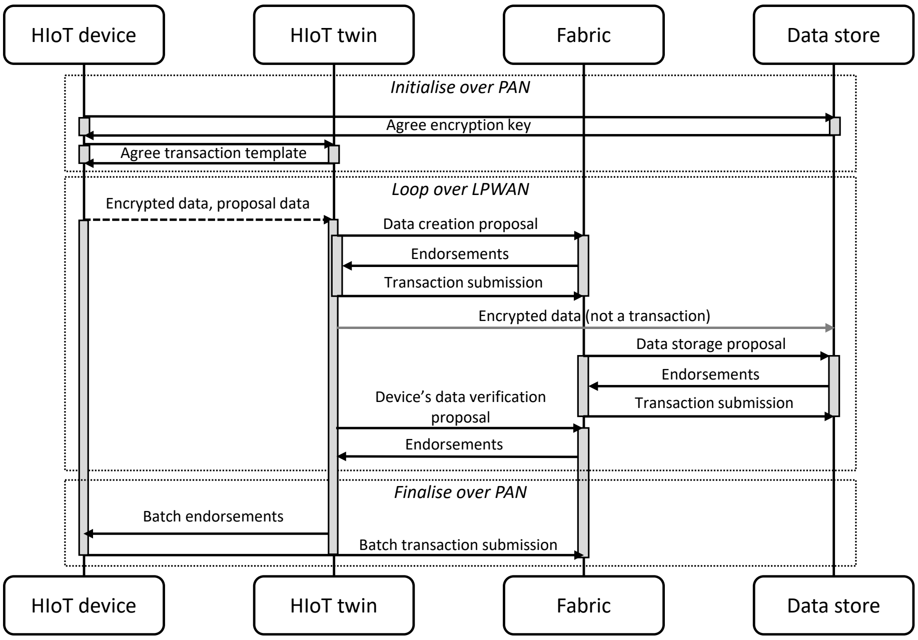

Figure 5. Sequence diagram of Fabric transactions representing HIoT sensor data early lifecycle.

The device and twin use the PAN to agree on a template for a proposal that verifies the data created by the device. Upon receiving a data packet over LPWAN, the twin can reconstruct this proposal for submission to the Fabric gateway. However, if it does so immediately, simulations of the proposal would fail, as it would refer to a non-existent data item.

Instead, the twin can propose its own transaction, using its own identity and private key, to execute chaincode that represents the creation of the data. The encrypted data are part of the payload they receive from the device, so they can produce a hash of it. The twin can also forward the same encrypted data to the data store, and the data store may have direct access to the data via the LPWAN's message queues (for example, an MQTT broker in the case of prominent LoRaWAN networks).

The data store, in possession of the encrypted data, should be able to decrypt them. It can also observe the twin's data creation transaction on the blockchain. Following this, it can submit a transaction that updates the status of this data item, indicating that it is intact and can be stored.

Observing the data store's transaction, the twin is now able to submit the device's proposal for endorsement. The endorsements can be collected, and once a PAN connection with the device is re-established, these can be relayed to the device for the creation of data validation transactions. The device's proposed transaction is created under the assumption that the other two transactions take place first. This is visible in the sequence diagram in the early activation of the device and twin at the start of the LPWAN loop portion, overlapping with two transactions by the twin and data store, before concluding with deactivation in the ending PAN communication portion. Multiple transactions may be batched together in this stage, as the LPWAN loop will have iterated many times between PAN-based synchronisations.

This approach more closely couples the existence of the data asset in the blockchain with the first transmission of it from the device, rather than deferring it until the next PAN connection. It also allows the data store to confirm the integrity and storage of the data before the device is finally able to confirm this also. It is a more fine-grained representation of the early stages of the data's lifecycle.

Figure 6 offers an alternative view of the same exchange. Data travel over LPWAN and PAN to the twin, which then interacts with both the Fabric gateway and the datastore, depicted by dashed arrows. In logical terms, this results in the device contributing data to the data store, along with the device, twin and data store contributing records to the blockchain that affect the PHR as depicted by dotted lines between said participants within Figure 6. The resulting data exchanges are numbered, with 1 being the sending of encrypted data from device to store, and 2, the record of the data's creation, which can be conceptualised as part of the patient's PHR on the blockchain. Subsequently, the data store can record its successful storage as item 3, and upon synchronisation over PAN, the final records of validity, 4, are entered into the blockchain to support the integrity of the PHR.

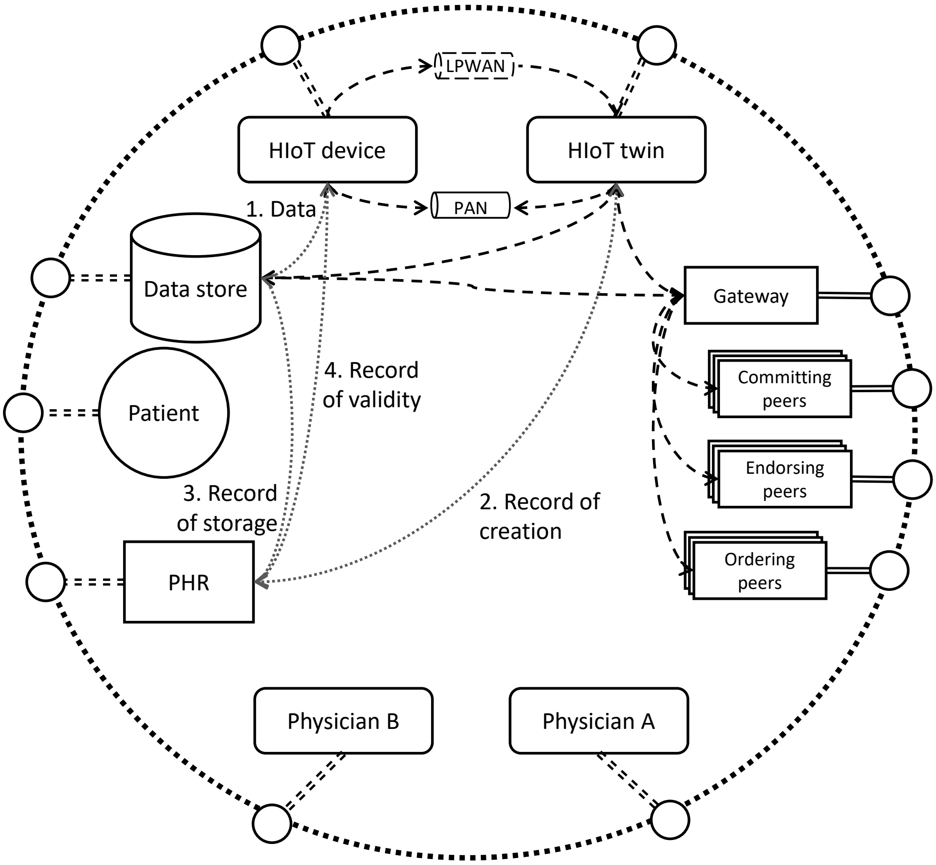

Figure 6. Paths of communication and resulting transactions in HIoT three-part early lifecycle. For legend, refer to Section 4.1.1.

4.2. Satisfying Security Requirements

In Section 3.4, we defined the security objectives sought by the HC^2 architecture in answer to our use case requirements from Section 2.4. Here, we explain how they are satisfied within the constraints of the technology choices made earlier within this section.

The primary security concerns and safeguards present in our system are summarised in Table 6 and explained in more detail in this section.

Table 6. CIA summary of HC^2. Some items discuss multiple security goals.

Confidentiality	Anonymity of data (Section 4.2.1), ownership of keys (Section 4.2.2), data keys (Section 4.2.3), off-chain data (Section 4.2.4), re-encryption (Section 4.2.5), forward secrecy (Section 4.2.7), post-quantum encryption (Section 4.2.8)
Integrity	Data keys (Section 4.2.3), post-quantum encryption (Section 4.2.8)
Availability	Missing data detection (Section 4.2.6)

4.2.1. Anonymity

No PII is transmitted by the device. Instead, the association between device and patient is maintained by reference in the blockchain. The identifiers used do not need to be directly attributable to a person; this can be resolved off-chain.

If device tracking is a concern, then additional countermeasures would be needed, such as changing device IDs and keys, for example, with each synchronisation. However, these are not considered further in this paper.

4.2.2. Device Key Ownership

The device can generate (or otherwise have injected) its own private key without being provided one by the Fabric CA, by leveraging the Fabric CA support for certificate signing requests (CSRs) during enrolment [37]. This precludes any possibility for impersonation of the device at the CA. If a TPM or a secure element is used on the device, the private key protection can be strengthened further [38].

Additionally, the twin does not share persistent key material with the device, so while both synchronise certain items of data (starting nonce, counter, and public keys), they cannot impersonate each other or tamper with signed messages. This remains the case despite the twin's ability to reconstruct signed messages from partial data transmitted from the device via LPWAN as described in Section 4.1.

4.2.3. Data Keys

During synchronisation between the device and twin, the twin also facilitates creating a secure session between the device and target data store. During this process, their respective identities are verified, and a symmetric encryption key is established for data transfer. The sensor data or events transmitted by the device are encrypted with this key, and thus the data store is the only other party able to decrypt it.

If the authentication encryption with associated data (AEAD) scheme, such as AES-GCM, is used, the encrypted data are accompanied by an authenticating tag that any party in possession of the symmetric key can use to verify the data integrity, independently of the message signature. In the case of 256-bit AES-GCM [39], the tag is 16 bytes, which must be included in the LPWAN transmission. Additionally, where AEAD is used, the integrity of the encrypted data is assured at this point, as well as later when the device verifies that it was stored.

4.2.4. Off-Chain Data

For privacy and efficiency, the sensor data are not stored on the blockchain. Instead, the hash of the encrypted data is stored. Any entity in possession of the encrypted data can verify that it is represented in the blockchain but can only decrypt them if in possession of the associated key.

Keeping data off-chain has the advantage of reducing the block sizes and growth rate of the blockchain by avoiding using the blockchain itself as a storage device. At significant scale, solutions such as IPFS may be used [4].

4.2.5. Re-Encryption of Data

The data store, or other accessors of the data, may re-encrypt the data to cease reliance on the key used between the device and data store. Provided the affected data assets can still be tracked, the integrity of the data in relation to the blockchain records can still be verified, provided the original encryption key is stored. This key should be stored with the equivalent protection as the re-encrypted data, for example, the original key could be stored encrypted by the new key.

4.2.6. Missing Data

Upon re-synchronisation, a device may additionally verify that the data it previously transmitted but also locally logged were indeed successfully stored. If they were not, a notification can be made. The data can then be provided during the synchronisation process instead. Although the benefits of real-time availability are lost, they will still eventually be available, and the evidence of their absence is provable.

4.2.7. Forward Secrecy

Forward secrecy is preserved through the use of ephemeral keys agreed between communicating parties. In the case of Fabric, this uses TLS. For the device and twin as well as device and data store, this may use DTLS [40] or EDHOC [41]. In all of these cases, the ephemeral keys used for data encryption are not related to the identifying keys of the participants. Thus, a successful attack on any of these ephemeral keys only affects data encrypted under that key. Each encrypted session must then be attacked independently.

In our use case, the session between the device and data store may last days or weeks, but the volume of data will not exceed a level that would pose a security risk through issues, such as initialisation vector reuse or exceeding data limits, which can affect AEAD ciphers, such as AES-GCM [39] (§8).

While we do not rely on the security of the LPWAN implementation for data or blockchain related activities, we remark that LoRaWAN agrees on a key during device activation [20] (pp. 62–63), [42] and that key management methods have been proposed or refined for it, too [43,44]. These could be more tightly integrated with blockchain identities and the Fabric CA/PKI if desired.

4.2.8. Post-Quantum Encryption

At the time of writing, post-quantum encryption (PQE) is a growing concern. Many of the cryptographic algorithms we use today are vulnerable to attack from the increased capabilities that quantum computers will eventually bring. New algorithms must be developed and adopted that are strong against conventional- and quantum-computing attacks but still feasible to run on conventional computers. For example, AES-256 encryption's security level is halved in the post-quantum area, and 256-bit EC-based key exchange and signing will be considered broken [45].

Institutions such as the USA's National Institute of Standards and Technology (NIST) continue to analyse and select candidate algorithms to address these concerns. However, these new algorithms are often more memory- and/or processor-intensive, which means they do not translate well to constrained IoT devices. Additionally, key and signature sizes in these PQE implementations can be significantly larger than those used today, making them unsuitable for use over LPWAN.

In early 2023 [46], NIST selected a family of lightweight cryptography algorithms targeting IoT and other constrained devices, named Ascon. These implement AEAD and hashing and so could substitute the existing algorithms that are part of the toolkits used in our demonstration codes. Additionally, one Ascon variant possesses some defences against quantum attacks; however, the NIST stance is that lightweight devices are less of a concern for PQE compared to systems responsible for long-term permanent storage.

In the case of our system, the data store may implement PQE and re-encrypt the data, using this to enhance protection. Forward secrecy remains in place on any data that were previously captured in transit for later decryption.

We do not explore the implications of PQE on algorithms used in the blockchain directly, as this is of interest to the community at large and not limited to HIoT. We do note, however, that the data are never stored in the blockchain, only a hash of the encrypted data (see Table 4).

5. Discussion and Limitations

In this section, we perform tests to analyse the efficiency of the HC^2 solution, discuss its performance and scaling properties, and consider the integration challenges faced when trying to develop an HC^2-enabled system.

5.1. *Scaling and Integration Considerations*

When deployed at scale, an HC^2 solution may encompass or interact with a variety of systems and many thousands of devices. In this section, we consider the efficiency of individual data packets, blockchain transactions, and the overall capacity of the blockchain, along with integration concerns. While we focus mainly on constraints relevant to regions that follow AS923 regulations, similar constraints must be considered in others, possibly with slightly differing results or optimal choices.

5.1.1. Data Payload Efficiency

At our proposed 125 byte LoRa payload (Table 5), 59 bytes are used by the encrypted sensor data, or 42.5% of the payload. If 16 bytes of that is also used for AEAD, 43 bytes of data remain, or 34.4%. In either case, more than half of the payload is used purely for blockchain-related data. If higher efficiency than this is required, then one must consider whether the benefits of the blockchain can be dispensed with, or substituted with a more lightweight alternative. Otherwise, a LoRa data rate that can accommodate a larger payload, or a different LPWAN technology altogether, may be preferred.

5.1.2. Fabric Payload Efficiency

This subsection examines the benefits brought by implementing the HC^2 scheme that we proposed when working within the transmission constraints of common LoRaWAN deployments. Our approach along with two alternatives are given as follows:

- **Hybrid Channel + Template**: The full HC^2 implementation, where we seek to send the bare-minimum non-templated data over LoRaWAN, relying on PAN, templates and the twin for data reconstruction and interactions with the Fabric gateway.
- **Hybrid Channel**: A simpler approach that still uses a PAN and twin to minimise LoRaWAN usage but uses a signed proposal generated and sent in full by the device.
- **Single Channel**: No PAN is used, and therefore messages for Fabric must be sent and received over LoRaWAN, even if a twin assists in transitioning between LoRaWAN and TCP/IP communication to the gateway.

An indicative set of data payloads is generated using our device and twin demonstration code [34] (Fabric samples: `scaling-data`), modified to output the length of the three messages that would be exchanged between the device and Fabric gateway (with or without twin assistance). They are the proposal, the signed endorsements for the proposal and finally the signed transaction. In the case of both hybrid channel variants, we assume that the endorsements and transaction are handled over PAN, which delays their processing but removes the need for LoRaWAN downlinks. However, for single channel, the downlinking of the endorsements would be required.

We chose a data size of 31 bytes, as this conveniently fit within our proof-of-concept use case whilst being a feasible length for the sensor data. It can be accommodated within a single LoRa packet in the Hybrid Channel + Template approach. When full Fabric messages are used, the message lengths vary due to the ASN.1 representation of signatures being 70–72 bytes long [47]. ECDSA signature values are represented as two signed integers, which may each need an additional octet to preserve their positive sign if the most-significant bit is set. This, combined with the headers in ASN.1, results in four possible lengths for an encoded signature. When determining the packet sizes in our demonstration code, we take the largest. This yields a proposal message of 1343 bytes, an endorsement message of 4709 bytes and a signed transaction of 4781 bytes.

Two LoRaWAN data rate profiles, DR5 and DR4, are used. In the higher data rate (DR5), a payload size of up to 222 bytes can be accommodated. However, HC^2 targets 125 byte payloads, which can be accommodated in both DR5 and DR4. We include both payload sizes and data rates in order to explore the effect that these choices have on the other transmission schemes that send full Fabric messages.

Figure 7 shows the number of LoRa packets needed when performing a single data transaction, that is, transmitting the encrypted sensor data, along with any signatures or other message components for Fabric, depending on the transmission scheme. HCT sets the baseline of using a single uplink packet, while HC uses more but benefits from the larger payload size available within DR5. The SC scheme, however, requires significantly more uplinks and also requires downlinks. Even with 222 byte payloads, the number of LoRa packets approaches fifty for each data transaction.

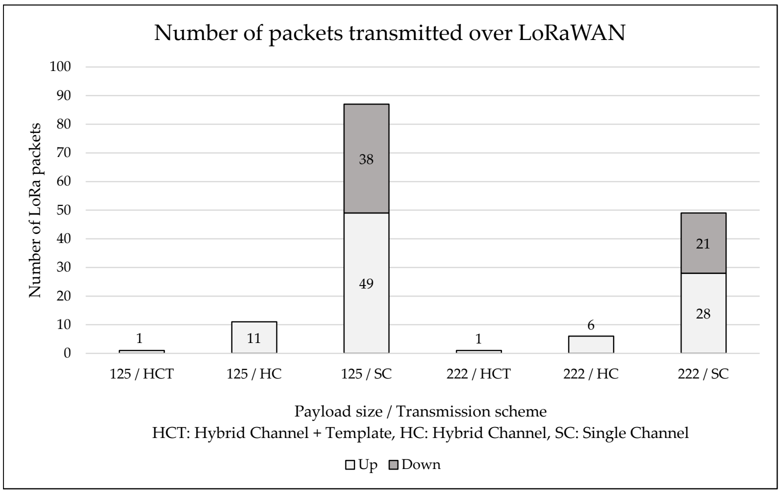

Figure 7. Packets transmitted in both directions (uplink and downlink) for each transmission scheme, using two available payload size limits over LoRa.

These data motivate the deferment of finalising transactions, as it significantly reduces LoRa utilisation (or allows data to be sent more frequently). If the three transactions proposed for HC^2 are used in the HCT and HC cases, this deferment is mitigated somewhat. With the assistance of templates, a further order of magnitude reduction in LoRa utilisation is achieved for 125 byte payloads. In the best case, HC^2 achieves an 87-times reduction in packets transmitted, thanks to combined message efficiencies and deferral.

Next, we examine the impact that duty cycle limits and fair access policies have on the amount of data that can be transmitted. As discussed in the literature view, regions impose limits on the amount of airtime that a LoRa device is able to occupy, in order to share the available bandwidth more fairly. Similarly, LoRaWAN providers may impose additional limits, such as even stricter duty cycles and limits on downlinks.

We take 1% as the duty cycle limit, which is applied in various global regions, including the AS923 region that is most relevant to the authors. To consider fair access policies, we use the things network (TTN), which limits airtime to 30 s per device per day and a maximum of ten downlinks. Figure 8 shows the results of applying these constraints to our selected payload sizes, data rates and transmission schemes.

To the left, Figure 8a shows that each of the three schemes are separated by an order of magnitude with respect to how many data transactions can be made per day. For HCT, the worst case is 2160 per day, or a data transaction every 40 s. HT is limited to every four minutes, whilst SC is limited to almost 33 min.

On the right, Figure 8b applies the TTN limits. Both HCT and HC are affected by the stricter airtime limits. SC incurs a much greater penalty due to the downlink limit, leading to two orders of magnitude, separating it from HC. In the 222-byte DR5 case, a small improvement is achieved due to more efficient use of the ten available downlinks per day because the larger transaction packets can be sent in fewer segments when the payload is larger; however, it does not substantially impact the results, as the transaction packets still exceed the payload size by several times.

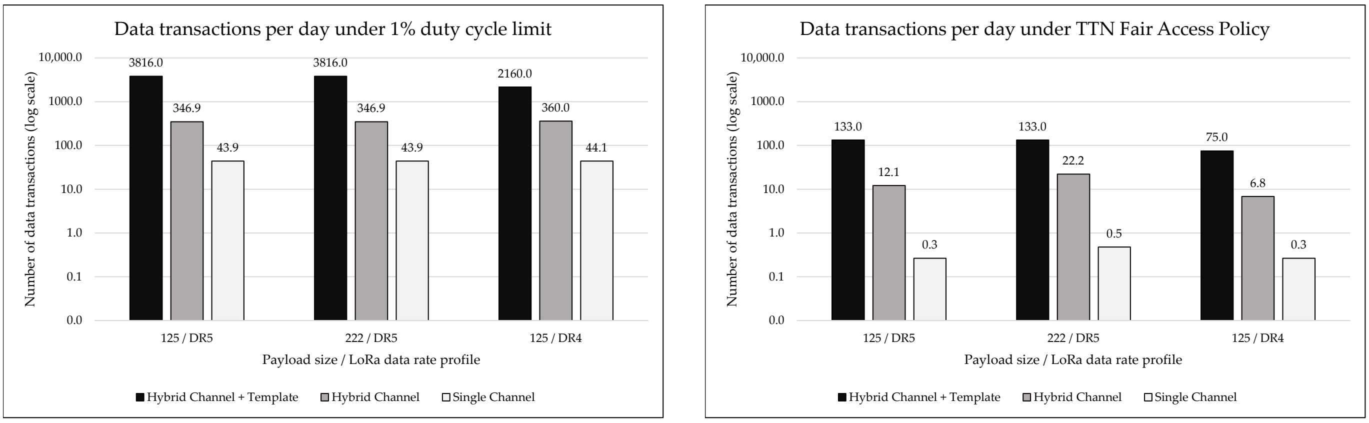

(**a**) With 1% duty cycle limit

(**b**) Under TTN limits, including maximum 10 downlinks

Figure 8. Maximum data transactions per day under selected transmission schemes, payload limits and LoRa data rate profiles. Scales are consistent between sub-figures.

These data demonstrate the feasibility of conducting Fabric transactions over LoRaWAN, making a case for avoiding downlinks where possible. Additionally, by utilising the HC^2 scheme, a further improvement can be obtained. Looking beyond just the data transfer limitations, there are also likely to be significant energy savings to be had for the device. Assuming a data transmission period is chosen that is not close to the limit, the full HC^2 limitation will use far less radio airtime, preserving battery life and potentially allowing for more advanced computation to be performed on the device with the spare energy.

5.1.3. Blockchain Utilisation

Using 3000 transactions per second (tps) as the baseline performance of Hyperledger Fabric [48], and the DR5 data rate (spreading factor 7 with 125 kHz bandwidth) combined with the TTN usage policy yielding 5.5 messages per hour per device [21,22], we calculate that over 650 thousand devices could be supported by the solution in terms of blockchain throughput, assuming each data transfer produces three transactions. Hyperledger Fabric can be scaled to higher transaction throughputs than this [48], although we speculate that any particular healthcare ecosystem on a single blockchain is not likely to exceed one million active HIoT devices. Various scaling enhancements, such as side chains, can be employed [49] should they be necessary and can be implemented with existing frameworks, including Fabric.

Another scaling limitation is the number of devices supported by each LoRa concentrator. This is affected by the amount of airtime each device's transmissions will use as well as the number of available channels, which is governed by concentrator support and regional regulations. Continuing to use the DR5 data rate, eight uplink channels is a moderate selection that can be accommodated by most regions and concentrators.

Equation (2) is a simple equation to determine the number of devices D, that can be accommodated on a channel, given a packet airtime A and a periodicity of transmission P for each device:

$$D = \left\lfloor \frac{P}{A} \right\rfloor \tag{2}$$

In the AS923 region, the DR5 date rate requires $A = 225.5\,\text{ms} = 0.2255\,\text{s}$ of airtime for each uplink packet of the 125 byte payload size used for HC^2 throughout this section. If devices each transmit at a five-minute period, $P = 300\,\text{s}$, then applying Equation (2), we find an ideal upper limit of $D = 1330$ devices that can be accommodated by a single concentrator. However, in a rural setting, the device density is likely to be lower, negating this concern. A deployment of 488 concentrators at the full density of 1330 devices per concentrator (with no overlap in reception) would be needed to approach the Fabric blockchain transaction limit.

While the three-part early lifecycle approach (Section 4.1.2) increases transactions on the blockchain by 3× versus a single deferred transaction, the dispersal of endorsements helps to reduce bottlenecks in the system. At synchronisation time, endorsements have already been collected by the twin for the device to batch together, reducing the number of transaction submission messages to Fabric. Provided the synchronisation of all participating HIoT devices is not performed at the same time, excessive load should be avoidable.

5.2. Integration Challenges

Blockchain technology remains an active area of research and development, and so future changes to blockchain technologies may create new challenges for integration into architectures, such as HC^2. In the case of the selections made in Table 1, we note two integration challenges that are avoided based on the present state of Hyperledger Fabric.

The first such challenge is the introduction of the *epoch* value into transactions. This numerical value represents the height of the block into which the transaction will go (i.e., the number of blocks in blockchain). If a transaction's epoch value is lower than the current height, the orderer will not include it. While a field in proposals and transactions is defined for this, presently, it is set to the value zero and thus is not enforced. Hyperledger Fabric JIRA issue FAB-1430 https://jira.hyperledger.org/browse/FAB-1430 (accessed on 27 April 2023) proposes checking of this epoch value, however the status of the work is "won't do" and the issue is closed. Therefore, at the time of writing, this feature is not expected to be implemented by the orderer, but given that provision exists within the message framework, if that decision is reversed, it would have negative implications for the proposed HC^2 implementation, as the IoT device would have way to track the current block height.

Secondly, Fabric supports mutual TLS, where the secure connection between the client and peer (i.e., twin and gateway) verifies both participants' identities. While this provides additional security in some contexts, without addition considerations, its use may prevent the twin from submitting proposals on behalf of the device, as the identity bound to the TLS channel would not match the identity of the submitted message. Additional application logic in Fabric and/or extensions within the issued certificates that associate device and twin with each other could overcome this without compromising the intent of mutual TLS.

Looking at the security integration between device and data store, using AEAD over the symmetric encryption session affects the data payload efficiency as discussed in Section 5.1.1. Under some circumstances it may be desirable to remove this, relying instead upon verification of the data payload by checking the signatures of blockchain messages from the device and twin. However, doing so requires tighter integration between the logic used in the blockchain and the logic of the storage application, which may not be desirable. We see this as a trade-off for which the decision may vary depending on use case and constraints.

The transfer of data through the HIoT system needs to be compatible with the integration with HC^2. Section 3 does not define any strict underlying requirements in the HIoT system, and Section 4 provides an example implementation that is refined based on the combination of LoRaWAN and Hyperledger Fabric. The messaging patterns, both in the model and implementation, may benefit from representation in a clearer form, such as that proposed in [50]. For example, TTN provides an MQTT data API for its LoRaWAN network, which can already represented with the «`MQTT`» stereotype in [50], and a similar stereotype may be created for Fabric gateway interactions. The templating implemented in HC^2 between device and twin can be formalised with «`ContentEnricher`» and «`EnvelopeWrapper`» to represent the transformation of messages as they transit between the device/twin and, subsequently, Fabric gateway. The different messaging formats and delivery methods over WAN and PAN should also be well-defined. A full UML definition of these patterns, or suitable equivalent, is beyond the scope of this paper, however.

6. Conclusions

In this paper, we made the case for tightly integrating HIoT data with the blockchain. This benefits the patient and healthcare provider by ensuring data integrity, increasing trust between parties. We also showed how the data can be transacted and stored without undue risk to patient confidentiality, such as the disclosure of PII or unencrypted storage/transit of sensor data.

We focused on how this integration can be delivered in rural settings, where network connectivity may be limited for potential patients but for whom the benefits of HIoT devices are still sought. A combination of communication channels, LPWAN and PAN, into hybrid-channel connectivity, along with a novel use of a digital twin and transaction templates, enables blockchain participation without overburdening devices or networks.

Our results show that with an appropriate implementation of the HC^2 model, blockchain-backed data transactions become feasible where they would not otherwise be, such as using LoRaWAN in combination with the Hyperledger Fabric blockchain. An 87x reduction in the number of LoRa transmissions needed is shown, allowing two orders of magnitude more data to be transferred under normal LoRaWAN operating constraints. Proof-of-concept code is provided that can serve as the basis for a full implementation, or as a comparison point for alternative solution proposals.

Future Work

For future work, we envisage two main pursuits. Firstly, a full implementation of the platform described in Section 4 to validate its effectiveness and explore its performance under real-world usage. Secondly, the exploration of alternative implementations of the HC^2 model, such as by using a different LPWAN technology or another blockchain system. Both of these areas of work would help to increase the understanding of the practical applications of our model and technical decisions that can maximise its benefits. Additionally, formalising the messaging patterns present in the hybrid-channel approach of HC^2, both for models and any implementations, may aid in integration efforts.

Author Contributions: Conceptualization, J.J., S.K. and T.H.; Methodology, J.J., S.K. and T.H.; Formal analysis, J.J. and S.K.; Writing—original draft preparation, J.J. and S.K.; Writing—review and editing, J.J., S.K. and T.H.; Funding acquisition, J.J., S.K. and T.H. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the competitive James Cook University Singapore—Australia Cross-Campus Collaboration Grant Scheme 2022.

Data Availability Statement: The data analysed in this work are generated from modifications to the Hyperledger Fabric SDK and application source code. These modifications, including the means to generate the data, are made available via a group of GitLab projects [34]. No other data collection exercises were performed.

Acknowledgments: The authors give thanks to Belinda Lee and Kevin Wang for their reserach administration support. We also wish to thank the reviewers for their time and feedback.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Qadri, Y.A.; Nauman, A.; Zikria, Y.B.; Vasilakos, A.V.; Kim, S.W. The Future of Healthcare Internet of Things: A Survey of Emerging Technologies. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 1121–1167. [CrossRef]
2. Li, J.; Carayon, P. Health Care 4.0: A vision for smart and connected health care. *IISE Trans. Healthc. Syst. Eng.* **2021**, *11*, 171–180. [CrossRef] [PubMed]
3. Dimitrievski, A.; Filiposka, S.; Melero, F.J.; Zdravevski, E.; Lameski, P.; Pires, I.M.; Garcia, N.M.; Lousado, J.P.; Trajkovik, V. Rural healthcare IoT architecture based on low-energy LoRa. *Int. J. Environ. Res. Public Health* **2021**, *18*, 7660. [CrossRef] [PubMed]
4. Dammak, B.; Turki, M.; Cheikhrouhou, S.; Baklouti, M.; Mars, R.; Dhahbi, A. LoRaChainCare: An IoT Architecture Integrating Blockchain and LoRa Network for Personal Health Care Data Monitoring. *Sensors* **2022**, *22*, 1497. [CrossRef] [PubMed]
5. Citoni, B.; Ansari, S.; Abbasi, Q.H.; Imran, M.A.; Hussain, S. Comparative Analysis of an Urban LoRaWAN Deployment: Real World Versus Simulation. *IEEE Sen. J.* **2022**, *22*, 17216–17223. [CrossRef]

6. Sun, Z.; Yang, H.; Liu, K.; Yin, Z.; Li, Z.; Xu, W. Recent Advances in LoRa: A Comprehensive Survey. *ACM Trans. Sens. Netw.* **2022**. [CrossRef]
7. Hathaliya, J.J.; Tanwar, S. An exhaustive survey on security and privacy issues in Healthcare 4.0. *Comput. Commun.* **2020**, *153*, 311–335. [CrossRef]
8. Chenthara, S.; Ahmed, K.; Wang, H.; Whittaker, F. Security and Privacy-Preserving Challenges of e-Health Solutions in Cloud Computing. *IEEE Access* **2019**, *7*, 74361–74382. [CrossRef]
9. Sun, P. Security and privacy protection in cloud computing: Discussions and challenges. *J. Netw. Comput. Appl.* **2020**, *160*, 102642. [CrossRef]
10. Pennino, D.; Pizzonia, M.; Vitaletti, A.; Zecchini, M. Blockchain as IoT Economy Enabler: A Review of Architectural Aspects. *J. Sens. Actuator Netw.* **2022**, *11*, 20. [CrossRef]
11. Wang, Y.; Su, Z.; Zhang, N. BSIS: Blockchain-Based Secure Incentive Scheme for Energy Delivery in Vehicular Energy Network. *IEEE Trans. Ind. Inform.* **2019**, *15*, 3620–3631. [CrossRef]
12. Farooq, M.S.; Suhail, M.; Qureshi, J.N.; Rustam, F.; de la Torre Díez, I.; Mazón, J.L.V.; Rodríguez, C.L.; Ashraf, I. Consortium Framework Using Blockchain for Asthma Healthcare in Pandemics. *Sensors* **2022**, *22*, 8582. [CrossRef]
13. Marbouh, D.; Simsekler, M.C.E.; Salah, K.; Jayaraman, R.; Ellahham, S. A Blockchain-Based Regulatory Framework for mHealth. *Data* **2022**, *7*, 177. [CrossRef]
14. Semwal, N.; Mukherjee, M.; Raj, C.; Arif, W. An IoT based smart e-health care system. *J. Inf. Optim. Sci.* **2019**, *40*, 1787–1800. [CrossRef]
15. Mekki, K.; Bajic, E.; Chaxel, F.; Meyer, F. A comparative study of LPWAN technologies for large-scale IoT deployment. *ICT Express* **2019**, *5*, 1–7. [CrossRef]
16. Xue, H.; Chen, D.; Zhang, N.; Dai, H.N.; Yu, K. Integration of blockchain and edge computing in internet of things: A survey. *Future Gener. Comput. Syst.* **2023**, *144*, 307–326. [CrossRef]
17. Glaessgen, E.; Stargel, D. The digital twin paradigm for future NASA and US Air Force vehicles. In Proceedings of the 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA, Honolulu, HI, USA, 23–26 April 2012; p. 1818. [CrossRef]
18. Liu, Y.; Zhang, L.; Yang, Y.; Zhou, L.; Ren, L.; Wang, F.; Liu, R.; Pang, Z.; Deen, M.J. A Novel Cloud-Based Framework for the Elderly Healthcare Services Using Digital Twin. *IEEE Access* **2019**, *7*, 49088–49101. [CrossRef]
19. Sahal, R.; Alsamhi, S.H.; Brown, K.N.; O'Shea, D.; McCarthy, C.; Guizani, M. Blockchain-Empowered Digital Twins Collaboration: Smart Transportation Use Case. *Machines* **2021**, *9*, 193. [CrossRef]
20. Sornin, N.; Yegin, A. *LoRaWAN™ 1.1 Specification 2*; Technical Report; LoRa Alliance, Inc.: Fremont, CA, USA, 2017.
21. van Bentem, A. Airtime Calculator for LoRaWAN. Available online: https://avbentem.github.io/airtime-calculator/ttn/as923/125 (accessed on 17 March 2023).
22. The Things Network. LoRaWAN Duty Cycle-Fair Use Policy. Available online: https://www.thethingsnetwork.org/docs/lorawan/duty-cycle/#fair-use-policy (accessed on 17 March 2023).
23. Baker, S.B.; Xiang, W.; Atkinson, I. Internet of Things for Smart Healthcare: Technologies, Challenges, and Opportunities. *IEEE Access* **2017**, *5*, 26521–26544. [CrossRef]
24. YIN, Y.; Zeng, Y.; Chen, X.; Fan, Y. The internet of things in healthcare: An overview. *J. Ind. Inf. Integr.* **2016**, *1*, 3–13. [CrossRef]
25. Gordon, W.J.; Catalini, C. Blockchain Technology for Healthcare: Facilitating the Transition to Patient-Driven Interoperability. *Comput. Struct. Biotechnol. J.* **2018**, *16*, 224–230. [CrossRef] [PubMed]
26. Haleem, A.; Javaid, M.; Singh, R.P.; Suman, R.; Rab, S. Blockchain technology applications in healthcare: An overview. *Int. J. Intell. Netw.* **2021**, *2*, 130–139. [CrossRef]
27. Maftei, A.A.; Mutescu, P.M.; Popa, V.; Petrariu, A.I.; Lavric, A. Internet of Things Healthcare Application: A Blockchain and LoRa Approach. In Proceedings of the 2021 International Conference on e-Health and Bioengineering (EHB), Iasi, Romania, 18–19 November 2021; pp. 1–4. [CrossRef]
28. Ahmad, R.W.; Salah, K.; Jayaraman, R.; Yaqoob, I.; Ellahham, S.; Omar, M. The role of blockchain technology in telehealth and telemedicine. *Int. J. Med. Inform.* **2021**, *148*, 104399. [CrossRef] [PubMed]
29. Adere, E.M. Blockchain in healthcare and IoT: A systematic literature review. *Array* **2022**, *14*, 100139. [CrossRef]
30. Munagala, N.V.L.M.K.; Rani, A.D.; Reddy, D.V.R.K. Blockchain-Based Internet-of-Things for Secure Transmission of Medical Data in Rural Areas. *Comput. J.* **2022**. [CrossRef]
31. Patel, V. A framework for secure and decentralized sharing of medical imaging data via blockchain consensus. *Health Inform. J.* **2019**, *25*, 1398–1411. [CrossRef]
32. Li, C.T.; Weng, C.Y.; Lee, C.C.; Wang, C.C. A Hash Based Remote User Authentication and Authenticated Key Agreement Scheme for the Integrated EPR Information System. *J. Med. Syst.* **2015**, *39*, 144. [CrossRef]
33. Agencia española de protección de datos. *Introduction to the Hash Function as a Personal Data Pseudonymisation Technique*; Technical Report; European Data Protection Supervisor: Brussels, Belgium, 2019.
34. Kerrison, S. HIoT Blockchain. Available online: https://gitlab.com/hiot-blockchain (accessed on 23 March 2023).
35. Hyperledger. Hyperleder Fabric: Key Concepts. Available online: https://hyperledger-fabric.readthedocs.io/en/release-2.4/key_concepts.html (accessed on 20 March 2023).

36. Sykes, M.; Yellick, J.; Enyeart, D. Hyperledger Fabric gRPC Service Definitions-Proposal. Available online: https://github.com/hyperledger/fabric-protos/blob/f0d57a53cb997351d8066fd6ab24cb48da1155b2/peer/proposal.proto (accessed on 17 March 2023).
37. Hyperledger. Hyperledger Fabric SDK for node.js. Available online: https://hyperledger.github.io/fabric-sdk-node/release-1.4/FabricCAClient.html#enroll__anchor (accessed on 20 March 2023).
38. Bouazzouni, M.A.; Conchon, E.; Peyrard, F. Trusted mobile computing: An overview of existing solutions. *Future Gener. Comput. Syst.* **2018**, *80*, 596–612. [CrossRef]
39. Dworkin, M.J. *SP 800-38D. Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC*; Technical Report; National Institute of Standards & Technology: Gaithersburg, MD, USA, 2007. [CrossRef]
40. Rescorla, E.; Tschofenig, H.; Modadugu, N. The Datagram Transport Layer Security (DTLS) Protocol Version 1.3. Technical Report; RFC Editor. April 2022. Available online: https://datatracker.ietf.org/doc/html/draft-carpenter-rfc-citation-recs-01#section-5.2 (accessed on 27 April 2023).
41. Selander, G.; Mattsson, J.P.; Palombini, F. *Ephemeral Diffie-Hellman Over COSE (EDHOC)*; Internet-Draft draft-ietf-lake-edhoc-19; Internet Engineering Task Force: Fremont, CA, USA, 2023; *Work in Progress*.
42. The Things Network. LoRaWAN Security. Available online: https://www.thethingsnetwork.org/docs/lorawan/security/ (accessed on 17 March 2023).
43. Sanchez-Iborra, R.; Sánchez-Gómez, J.; Pérez, S.; Fernández, P.J.; Santa, J.; Hernández-Ramos, J.L.; Skarmeta, A.F. Enhancing LoRaWAN Security through a Lightweight and Authenticated Key Management Approach. *Sensors* **2018**, *18*, 1833. [CrossRef]
44. Chen, X.; Lech, M.; Wang, L. A Complete Key Management Scheme for LoRaWAN v1.1. *Sensors* **2021**, *21*, 2962. [CrossRef]
45. Bernstein, D.; Lange, T. Post-quantum cryptography. *Nature* **2017**, *549*, 188–194. [CrossRef]
46. Boutin, C. NIST Selects 'Lightweight Cryptography' Algorithms to Protect Small Devices. Available online: https://www.nist.gov/news-events/news/2023/02/nist-selects-lightweight-cryptography-algorithms-protect-small-devices (accessed on 20 March 2023).
47. Housley, R.; Polk, T.; Bassham, L. Algorithms and Identifiers for the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile; Technical Report; RFC Editor. April 2002. Available online: https://datatracker.ietf.org/doc/html/draft-carpenter-rfc-citation-recs-01#section-5.2 (accessed on 27 April 2023).
48. Gorenflo, C.; Lee, S.; Golab, L.; Keshav, S. FastFabric: Scaling Hyperledger Fabric to 20,000 Transactions per Second. In Proceedings of the 2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), Seoul, Republic of Korea, 14–17 May 2019; pp. 455–463. [CrossRef]
49. Kim, S.; Kwon, Y.; Cho, S. A Survey of Scalability Solutions on Blockchain. In Proceedings of the 2018 International Conference on Information and Communication Technology Convergence (ICTC), Jeju, Republic of Korea, 17–19 October 2018; pp. 1204–1207. [CrossRef]
50. Górski, T. UML Profile for Messaging Patterns in Service-Oriented Architecture, Microservices, and Internet of Things. *Appl. Sci.* **2022**, *12*, 12790. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.