



UNIVERSIDADE FEDERAL DE SANTA CATARINA
CENTRO TECNOLÓGICO
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

Daniel Dagostim

**Enhancing Federated Learning in Massive IoT via Contention Resolution
Diversity Slotted ALOHA**

Florianópolis
2026

Daniel Dagostim

**Enhancing Federated Learning in Massive IoT via Contention Resolution
Diversity Slotted ALOHA**

Dissertação submetida ao Programa de Pós-Graduação
em Engenharia Elétrica da Universidade Federal de
Santa Catarina para a obtenção do título de mestre
em Engenharia Elétrica.

Orientador: Prof. Richard Demo Souza, Dr.

Coorientador: Prof. Jamil Farhat, Dr.

Florianópolis

2026

Ficha catalográfica gerada por meio de sistema automatizado gerenciado pela BU/UFSC.
Dados inseridos pelo próprio autor.

Dagostim, Daniel
Enhancing Federated Learning in Massive IoT via
Contention Resolution Diversity Slotted ALOHA / Daniel
Dagostim ; orientador, Richard Demo Souza, coorientador,
Jamil Farhat, 2026.
47 p.

Dissertação (mestrado) - Universidade Federal de Santa
Catarina, Centro Tecnológico, Programa de Pós-Graduação em
Engenharia Elétrica, Florianópolis, 2026.

Inclui referências.

1. Engenharia Elétrica. 2. Aprendizado Federado. 3.
Protocolo de acesso aleatório. 4. CRDSA. I. Demo Souza,
Richard. II. Farhat, Jamil. III. Universidade Federal de
Santa Catarina. Programa de Pós-Graduação em Engenharia
Elétrica. IV. Título.

Daniel Dagostim

**Enhancing Federated Learning in Massive IoT via Contention Resolution
Diversity Slotted ALOHA**

O presente trabalho em nível de mestrado foi avaliado e aprovado por banca examinadora composta pelos seguintes membros:

Prof. Ana Flávia dos Reis, Dr.
Universidade Federal do Paraná

Prof. Glauber Gomes de Oliveira Brante, Dr.
Universidade Tecnológica Federal do Paraná

Certificamos que esta é a **versão original e final** do trabalho de conclusão que foi julgado adequado para obtenção do título de mestre em Engenharia Elétrica.



Documento assinado digitalmente
Carlos Renato Rambo
Data: 09/03/2026 12:47:59-0300
CPF: ***.748.198-**
Verifique as assinaturas em <https://v.ufsc.br>

Prof. Carlos Renato Rambo, Dr.
Coordenador do Programa



Documento assinado digitalmente
Richard Demo Souza
Data: 06/03/2026 11:37:34-0300
CPF: ***.267.379-**
Verifique as assinaturas em <https://v.ufsc.br>

Prof. Richard Demo Souza, Dr.
Orientador

Florianópolis, 2026.

AGRADECIMENTOS

Em primeiro lugar, agradeço a Deus, que me permitiu viver e brindar à vida, explorando meus potenciais. Também agradeço e dedico este projeto aos meus pais, à minha família e aos meus amigos, pelo suporte ao longo de toda a minha trajetória desde cedo, em especial àqueles que galgaram o caminho antes de mim e foram fontes de inspiração, pois, como disse Newton, "se vi mais longe, foi por estar sobre os ombros de gigantes".

Agradeço à Intelbras S/A, empresa que tem como política fomentar o desenvolvimento de seus colaboradores e que, portanto, me concedeu abono de horas em algumas disciplinas cursadas durante o mestrado. Aos colegas e às lideranças, meu muito obrigado pelo apoio.

Agradeço ao PPGEEL da UFSC por proporcionar opções de pesquisa em distintas áreas, bem como por aprofundar e expandir o conhecimento por meio da ciência aplicada e de tecnologias de inovação para toda a comunidade.

Agradeço aos professores, em especial ao meu orientador Richard Souza e ao meu coorientador Jamil Fahat, pelos conselhos, pelas sugestões, pelo tempo investido e, principalmente, pela paciência e dedicação dispensadas a este trabalho, o que certamente foi fundamental em todo o processo.

Por fim, agradeço à minha esposa, Amanda, que muito me apoiou ao longo de todos os anos de pós-graduação e me compreendeu nas ausências inevitáveis. Este trabalho também é por ela.

Este trabalho foi parcialmente apoiado pelo CNPq/INCT STREAM (Signal Processing and Transmission for Environmental Analysis and Monitoring) - 409179/2024-8.

"The three great essentials to achieve anything worthwhile are, first, hard work; second, stick-to-itiveness; third, common sense."
Thomas A. Edison

RESUMO

O aprendizado federado é um paradigma promissor para o treinamento colaborativo de modelos de aprendizado de máquina em dispositivos descentralizados. Ao evitar o compartilhamento de dados brutos e sensíveis, este paradigma preserva a privacidade dos usuários e a integridade do sistema. Recentemente, o sucesso dessa abordagem em redes de telecomunicações impulsionou sua aplicação na Internet das Coisas, ambiente onde a segurança de dados é crítica, mas cujos dispositivos possuem restrições severas de processamento e energia. Neste contexto, o desempenho de sistemas de aprendizado federado é sensível à eficiência dos mecanismos de acesso aleatório no canal de comunicação. Protocolos tradicionais, como o *Slotted ALOHA* (SA), sofrem com colisões frequentes que retardam a convergência do modelo e limitam a escalabilidade da rede. Este trabalho investiga o uso de *Contention Resolution Diversity Slotted ALOHA* (CRDSA) em cenários de uso massivo de Internet das Coisas. A partir de um mecanismo de controle de acesso baseado em *feedback*, onde a probabilidade de acesso depende dessa realimentação que estabelece o quão congestionados os canais estão, e com base nisso, adaptamos o CRDSA para maximizar a taxa de sucesso de transmissão (*throughput*) e reduzir a perda de pacotes durante a etapa de agregação. Resultados de simulação demonstram que o CRDSA acelera significativamente a convergência e reduz a norma do erro em comparação ao SA e ao Polling, especialmente em cenários com disponibilidade suficiente de recursos de rádio. Os resultados indicam que o CRDSA é uma solução escalável e eficaz para redes de Internet das Coisas habilitadas com aprendizado federado.

Palavras-chave: Aprendizado Federado, Internet das Coisas, ALOHA.

ABSTRACT

Federated Learning is a promising paradigm for the collaborative training of machine learning models across decentralized devices. By avoiding the exchange of raw and sensitive data, this paradigm preserves user privacy and system integrity. Recently, the success of this approach in telecommunications networks has driven its application in the Internet of Things, an environment where data security is critical, yet devices face severe processing and energy constraints. However, the performance of federated learning systems is highly sensitive to the efficiency of random access mechanisms in the communication channel. Traditional protocols, such as *Slotted ALOHA* (SA), suffer from frequent collisions that slow down model convergence and limit network scalability. This work investigates the use of *Contention Resolution Diversity Slotted ALOHA* (CRDSA) in massive Internet of Things scenarios. Leveraging a feedback-based access control mechanism, we adapt CRDSA to maximize throughput and reduce packet loss during the model aggregation phase. Simulation results demonstrate that CRDSA significantly accelerates convergence and reduces the error norm compared to SA and Polling, particularly when sufficient radio resources are available. These findings indicate that CRDSA is a scalable and effective access solution for Federated Learning-enabled Internet of Things networks.

Keywords: Federated Learning, Internet of Things, ALOHA.

RESUMO EXPANDIDO

Introdução

Em 2025, aproximadamente 19,8 bilhões de dispositivos de Internet das Coisas (*Internet of Things*, IoT) estavam conectados globalmente, com projeções superando a marca de 30 bilhões até 2030. Os cenários atuais de IoT, caracterizados pela densidade massiva de dispositivos e volumes de dados sem precedentes, constituem um ambiente propício para o desenvolvimento de diversas aplicações. Consequentemente, tornam-se essenciais novas abordagens de processamento de dados que atendam a requisitos rigorosos de comunicação rápida, eficiente e ininterrupta, sem negligenciar aspectos críticos de segurança e privacidade.

Neste contexto, o aprendizado de máquina tem sido progressivamente integrado aos sistemas IoT. Entretanto, privacidade de dados permanece como um desafio central. Enquanto os ambientes IoT geram e trocam vastas quantidades de informações continuamente, as abordagens tradicionais e centralizadas, que agregam dados brutos em um *data center* para treinamento, apresentam limitações tanto de privacidade quanto de escalabilidade. Para mitigar esses problemas, o aprendizado federado emergiu como um paradigma promissor. No aprendizado federado, o treinamento do modelo é realizado localmente nos dispositivos terminais utilizando dados privados; em contrapartida, apenas os gradientes ou atualizações do modelo são compartilhados com uma estação base para agregação global. Este processo evita a transmissão de dados brutos sensíveis, o que não apenas fortalece a proteção da privacidade, mas também potencialmente reduz o custo de comunicação associado ao envio de grandes conjuntos de dados. Além disso, a otimização federada desempenha um papel fundamental na eficiência do sistema, especialmente quando modelos são treinados sobre dados altamente distribuídos e sob ciclos de transmissão frequentes. Ao combinar técnicas de aprendizado de máquina com IoT, o aprendizado federado viabiliza o processamento em larga escala e melhora a generalização do sistema, uma vez que os modelos locais capturam características específicas dos usuários antes da unificação global.

Para assegurar a eficiência diante da dependência de treinamentos locais e transmissões frequentes, a comunicação deve ser criteriosamente otimizada, particularmente no *uplink*. Esquemas sincronizados tradicionais, como *Time Division Multiple Access* (TDMA) e *Frequency Division Multiple Access* (FDMA), são inerentemente livres de colisão. Contudo, sua complexidade de implementação e custo operacional os tornam mais adequados a cenários de tráfego intenso ou missões críticas. Por outro lado, sistemas de aprendizado federado beneficiam-se de soluções leves e escaláveis que tolerem certo nível de colisão, desde que a convergência do modelo seja preservada. Por essa razão, protocolos de acesso aleatório *grant-free* (sem necessidade de con-

cessão) são particularmente atraentes, pois minimizam o *overhead* de sinalização e alinham-se à natureza dinâmica e descentralizada das redes IoT massivas.

Sob essa perspectiva, o desenvolvimento de estratégias de acesso aleatório eficazes torna-se um facilitador-chave para sistemas IoT baseados em aprendizado federado. Dentre as diversas técnicas, o ALOHA destaca-se como uma opção viável para o *uplink*, reduzindo custos e permitindo que dispositivos em modo de repouso transmitam atualizações sem concessão prévia. Por ser um protocolo baseado em contenção e *grant-free*, o ALOHA adequa-se ao tráfego esporádico da IoT, embora sua simplicidade resulte em limitações de *throughput* devido à elevada taxa de colisões em cenários densos.

Neste trabalho, propõe-se uma estratégia de comunicação para aplicações de aprendizado federado via protocolo *Contention Resolution Diversity Slotted ALOHA* (CRDSA). Esta abordagem é apresentada como uma alternativa superior ao *Polling* e ao *Slotted ALOHA* (SA) convencional. Por meio de uma análise comparativa, são realizadas simulações para validar a modelagem analítica desenvolvida, seguidas de uma discussão sobre a seleção de parâmetros ótimos baseada nas características intrínsecas do CRDSA. Os resultados demonstram que o número de iterações necessárias para atingir uma norma de erro reduzida, métrica fundamental de eficiência em sistemas federados, é mitigado em 50% e 85% quando comparado ao SA e ao *Polling*, respectivamente.

Objetivos

Esta dissertação tem como objetivo investigar a eficiência da integração entre o protocolo CRDSA e sistemas de aprendizado federado, analisando seu impacto no desempenho da rede e na convergência de modelos em cenários alto volume de dispositivos. Os seguintes objetivos específicos são: (i) desenvolver a modelagem para o comportamento do protocolo CRDSA sob as premissas e dinâmicas do aprendizado federado; (ii) analisar o impacto de parâmetros do sistema, tais como o número de canais ortogonais, o número de iterações e a probabilidade de computação local, no desempenho da comunicação; (iii) quantificar os ganhos de desempenho proporcionados pelo CRDSA em comparação ao protocolo SA e ao método de *Polling*.

Metodologia

Inicialmente, estabelece-se a base teórica sobre o aprendizado federado, explorando seus mecanismos de comunicação e os protocolos de acesso aleatório, com ênfase

no ALOHA e suas variações, SA e CRDSA. Realiza-se uma revisão bibliográfica para identificar o estado da arte em sistemas de aprendizado federado habilitados por protocolos de acesso aleatório. Este trabalho adota como base a arquitetura proposta por Choi et al. (2019), que investiga a comunicação de *uplink* para o envio de atualizações locais em sistemas de aprendizado federado. O referido modelo contrasta o desempenho do protocolo SA com o método livre de colisões *Polling*.

Resultados e discussão

A análise de resultados inicia-se pela modelagem da carga de tráfego e das expressões para a probabilidade de acesso, estabelecendo-se um comparativo fundamental entre o SA e o CRDSA. Na sequência, identifica-se a configuração ideal do CRDSA quanto ao número de réplicas, determinando-se a quantidade que maximiza o *throughput* em relação ao SA. Após a otimização desses parâmetros, é feita a análise da evolução da norma de erro em função do número de canais disponíveis. Na sequência, avaliou-se a convergência em função do número de iterações. Os resultados ratificam a superioridade do CRDSA, que atinge o mesmo patamar de norma de erro que o SA utilizando aproximadamente muito menos iterações.

Finalmente, é analisado o impacto da capacidade de processamento local, variando a probabilidade de computação sob um regime fixo de 100 iterações. O CRDSA apresenta o melhor desempenho até o limiar de probabilidade de computação de 30%. Acima deste valor, o SA torna-se ligeiramente superior e, ao ultrapassar 55%, o método de *Polling* supera ambos os protocolos de acesso aleatório. Esta transição ocorre porque o *Polling*, por ser livre de colisões e baseado em agendamento sequencial, torna-se mais eficiente quando há um volume muito alto de dispositivos prontos para transmitir simultaneamente. Contudo, em cenários típicos de aprendizado federado, as características dos dispositivos e dos ambientes de aplicação geralmente resultam em baixas probabilidades de computação, consolidando o CRDSA como a solução mais equilibrada e eficaz.

Considerações finais

Esta dissertação propôs a aplicação inédita do protocolo CRDSA em sistemas de aprendizado federado, preenchendo uma lacuna na literatura que, até então, focava majoritariamente no uso do *Slotted ALOHA*. Através da modelagem analítica e de simulações, foram determinados os valores ótimos de carga de tráfego, o número ideal

de réplicas e a probabilidade de acesso necessária para a correta operação do sistema sob as dinâmicas do aprendizado federado.

Os resultados obtidos confirmam a superioridade do método proposto. Esta eficácia é determinante em ecossistemas de aprendizado federado caracterizados pelo número massivo de dispositivos e transmissões esporádicas. A redução no tempo necessário para a convergência do modelo global valida o CRDSA como uma tecnologia facilitadora para a privacidade e eficiência em IoT. Como trabalho futuro, vislumbra-se a investigação do protocolo *Irregular Repetition Slotted ALOHA* (IRSA), que pode oferecer ganhos ainda maiores de *throughput* ao explorar a diversidade temporal de forma adaptativa, permitindo que as réplicas dos pacotes sejam distribuídas de maneira otimizada entre os slots de transmissão.

Palavras-chave: IoT. Acesso aleatório. CRDSA.

LIST OF FIGURES

Figure 1 – The iterative Federated Learning process within a single-cell network. The sequence includes the broadcast of the global model, local training at each device, and the transmission of local updates back to the base station for aggregation.	20
Figure 2 – Packet transmission and collision resolution in CRDSA. During iteration t , devices send updates to the BS. The SIC process iteratively cancels replicas of decoded packets, enabling further packet recovery.	22
Figure 3 – Polling Protocol: The BS polls devices sequentially. In this framework, the uplink transmission itself serves as an implicit acknowledgment; if a device is computationally active and polled, it transmits its update. If it is inactive, the assigned channel remains idle.	27
Figure 4 – Throughput versus normalized offered load for Slotted ALOHA and CRDSA with different replication factors $R \in \{2, 3, 4\}$	33
Figure 5 – Convergence of the access probability as a function of the number of iterations for CRDSA and SA. Raw results and their moving average are shown, where the moving average of the true aggregate access probability $\hat{P}_t$ is computed over a sliding window of the most recent t iterations, such that at each iteration the newest value of $\hat{P}_t$ is included and the oldest is discarded.	35
Figure 6 – Impact of the number of replicas on the error norm $\ w(t) - w^*\ $ at iteration 100.	36
Figure 7 – Error norm $\ w(t) - w^*\ $ at iteration 100 as a function of the number of available orthogonal channels, for CRDSA, Slotted ALOHA, and Polling.	37
Figure 8 – Convergence measured by the error norm $\ w(t) - w^*\ $ versus the number of iterations, comparing CRDSA, Slotted ALOHA, and Polling.	38
Figure 9 – Error norm $\ w(t) - w^*\ $ at iteration 100 as a function of the number of users, for CRDSA, Slotted ALOHA, and Polling.	39
Figure 10 – Error norm $\ w(t) - w^*\ $ at iteration 100 versus the probability p_{comp} at iteration 100.	40

LIST OF ACRONYMS

AoG	Age of Gradient
BS	Base Station
CRDSA	Contention Resolution Diversity Slotted ALOHA
D2D	Device to Device
DFL	Decentralized Federated Learning
DSA	Diversity Slotted ALOHA
DVB-RCS2	Digital Video Broadcasting – Return Channel via Satellite, 2nd Generation
FDMA	Frequency Division Multiple Access
FedAvg	Federated Averaging
FedPNS	Federated Probabilistic Node Selection
FL	Federated Learning
IEEE	Institute of Electrical and Electronics Engineers
IID	Independent and Identically Distributed
IoT	Internet of Things
IRSA	Irregular Repetition Slotted ALOHA
KKT	Karush-Kuhn-Tucker conditions
LoRaWAN	Long Range Wide Area Network
MAC	Medium Access Control
ML	Machine Learning
MNIST	Modified National Institute of Standards and Technology
non-IID	Non-Independent and Identically Distributed
RA	Random Access
SA	Slotted ALOHA
SGD	Stochastic Gradient Descent
SIC	Successive Interference Cancellation
TDMA	Time Division Multiple Access
WCNPS	Workshop on Communication Networks and Power Systems

LIST OF SYMBOLS

G	Traffic Load or Offered Load
T	Throughput
M	Number of channels
N	Total number of transmissions
R	Number of replicas
K	Number of Users
w^*	Optimal model parameters
D_k	Local Dataset
x_k	Input Samples data vector
y_k	Output samples data vector
w	Model parameters (weight vector)
p_{comp}	Probability of a packet successfully being computed
S_t	The set of devices which updates were successfully received
w_{t+1}	Model update
ρ_k	Access Probability
Q	Expected number of successfully received updates per iteration
δ_k	Indicator variable of successful reception from user k
λ	Lagrange multiplier
$\hat{P}_t$	Estimate of P at iteration t
$s_{k,t}$	Activity variable of user k at iteration t
μ_ψ	Step size (learning rate) of feedback update
G_R^*	The ideal traffic load of CRDSA
μ_w	Step size (learning rate) of local SGD update

CONTENTS

1	INTRODUCTION	16
1.1	PUBLICATION	18
2	BACKGROUND	19
2.1	FEDERATED LEARNING	19
2.2	RANDOM ACCESS FOR FEDERATED LEARNING	21
2.2.1	Slotted ALOHA	21
2.2.2	Contention Resolution Diversity Slotted ALOHA	22
2.3	RELATED WORK	23
2.3.1	Baseline Work	23
2.3.2	Alternative Approaches	23
2.3.3	Positioning of This Work	24
3	SYSTEM MODEL	26
3.1	FEDERATED LEARNING OVER POLLING	27
3.2	FEDERATED LEARNING OVER SLOTTED ALOHA	28
3.3	FEDERATED LEARNING OVER CRDSA	31
4	RESULTS	34
4.1	NUMERICAL RESULTS	34
5	CONCLUSIONS AND FUTURE WORK	41
	REFERENCES	43

1 INTRODUCTION

In 2025, approximately 19.8 billion IoT devices were connected worldwide, and this number is projected to exceed 30 billion by 2030 (Statista, 2025). Current IoT scenarios, characterized by an enormous number of devices and massive volumes of transmitted data, provide fertile ground for the development of diverse applications. Consequently, new approaches for data processing are essential, with stringent requirements for fast, efficient, and seamless communication. At the same time, critical aspects such as security and privacy must be carefully addressed (Nguyen et al., 2021; Adam et al., 2024).

In this context, Machine Learning (ML) is increasingly being integrated into IoT systems. However, data privacy remains a critical concern. IoT environments continuously generate and exchange vast amounts of data, while traditional centralized ML approaches, which aggregate raw data in a data center for training, pose both privacy and scalability challenges (Gao et al., 2022). To address this, Federated Learning (FL) has emerged as a promising paradigm (Elzemity et al., 2024). In FL, user devices perform local model training using their private data. In contrast, only the learned gradients or model updates are shared with the base station for aggregation into a global model (Zhang et al., 2025). This process avoids the exchange of sensitive raw data, which not only strengthens privacy protection but also may reduce the communication burden associated with transmitting large initial datasets. Furthermore, federated optimization plays a key role in ensuring communication efficiency when training centralized models with highly distributed data, as transmission cycles become more frequent compared to conventional ML (Gao et al., 2022). By combining ML techniques with IoT, FL enables large-scale data processing while enhancing privacy (Gao et al., 2022; Perazzone et al., 2022; Adam et al., 2024). In addition, it provides better system-level generalization, since local models capture user-specific data characteristics before being aggregated into a unified global model (Elzemity et al., 2024).

To maintain efficiency, given the reliance on locally trained data and frequent transmission cycles, communication must be carefully optimized, particularly in the uplink (Chen, M. et al., 2021). Traditional synchronized schemes, such as Time Division Multiple Access (TDMA) and Frequency Division Multiple Access (FDMA), are collision-free by design (Lun et al., 2024). However, their higher implementation complexity and cost make them suitable primarily for heavy-traffic scenarios or environments where collisions and packet losses must be strictly avoided (Chen, M. et al., 2021). Instead, FL scenarios benefit from lightweight, cost-effective solutions that tolerate a certain level of collisions while maintaining overall convergence performance (Valente da Silva et al., 2023). For this reason, grant-free random access protocols are particularly attractive, since they minimize signaling overhead, support massive numbers of devices, and align

naturally with the decentralized and dynamic nature of IoT networks (Silva, R. V. da et al., 2022).

From this perspective, the design of suitable random access strategies becomes a key enabler for FL-based IoT systems. Among different Random Access (RA) techniques, ALOHA stands out as a suitable option for uplink communication, as it reduces overhead and implementation costs while allowing devices, often dormant for long periods, to transmit updates without prior permission (Valente da Silva et al., 2023). Specifically, this strategy is a simple, grant-free, contention-based protocol that tolerates collisions, making it well aligned with sporadic IoT traffic (Huang et al., 2021). Its simplicity and asynchronous operation come at the expense of throughput, which is limited due to frequent collisions (Seo et al., 2025).

To improve the efficiency of RA schemes, Slotted Aloha (SA) was introduced, which considerably increases the overall throughput by synchronizing transmissions into well-defined time slots, thus reducing the probability of destructive collisions between packets (Seo et al., 2025; Dossa et al., 2024). Building upon this concept, further enhancements led to the development of Diversity Slotted ALOHA (DSA), where multiple replicas of the same packet are transmitted in different slots to exploit time diversity and improve the chances of at least one successful reception (Celandroni et al., 2009). While this approach increases reliability under moderate traffic conditions, under heavy traffic loads the excessive number of replicas significantly intensifies contention, causing a sharp increase in collisions and ultimately limiting the scalability of the system (Celandroni et al., 2009). To address this critical limitation, CRDSA was later proposed, leveraging the Successive Interference Cancellation (SIC) technique to iteratively decode and remove collided packets, thereby recovering more transmissions and boosting the achievable throughput even at higher channel loads (Kumazawa et al., 2023).

In practice, ALOHA-variants have been widely adopted in systems such as LoRaWAN, Sigfox, and DVB-RCS2, particularly in scenarios where many devices remain idle for extended periods (Dossa et al., 2024; Mikhaylov et al., 2020; Iida et al., 2022). Their lightweight, grant-free nature makes them well-suited to the communication patterns of massive IoT deployments and to the decentralized, iterative model updates of FL. In this context, slotted schemes are especially attractive: they preserve the simplicity of random access while improving efficiency through light synchronization. For this reason, SA-based methods have been advocated as practical solutions for the data link layer in massive IoT and FL applications (Choi et al., 2019).

Building on this foundation, the present work investigates the use of CRDSA to enable scalable and efficient communication for FL-based IoT systems, considering a SA baseline as reference. The remainder of this work is organized as follows: Chapter 2 provides some theoretical background and reviews related literature. Chapter 3 descri-

bes the system model and derives the analytical expressions governing the CRDSA implementation. In Chapter 4, we present and discuss both analytical and simulation results. Finally, Chapter 5 concludes the work and suggests future research directions.

1.1 PUBLICATION

The research discussed in this work led to the paper “Enhancing Federated Learning in Massive IoT via Contention Resolution Diversity Slotted ALOHA,” authored by D. Dagostim, J. Farhat, and R. D. Souza. It was presented at the **2025 Workshop on Communication Networks and Power Systems (WCNPS)** in Brasília, Brazil, where it was recognized with an *Honorable Mention* among the communication networks papers. The article has since been published on the IEEE Xplore Digital Library with the DOI: 10.1109/WCNPS69127.2025.11295873.

2 BACKGROUND

2.1 FEDERATED LEARNING

Federated Learning was introduced by McMahan et al. (2017) as a distributed machine learning paradigm that enables collaborative model training across decentralized devices without requiring the exchange of raw data. In FL, each device trains a local model using its private dataset and transmits only the resulting model update, typically a gradient or a weight vector, to a central server for aggregation. This approach preserves user privacy and reduces communication overhead compared to centralized learning schemes that require uploading large datasets.

McMahan et al. (2017), also developed a new algorithm coined Federated Averaging (FedAvg). This new algorithm enabled collaborative model training across decentralized devices and addressed two main challenges of typical federated optimization: non-independent and identically distributed data (non-IID) and unbalanced data distributions. In such settings, data samples differ among users, and clients may contribute unequally to the training process depending on the size and nature of their local datasets.

The FedAvg algorithm is derived from SGD, a widely used optimization method for training ML models. Instead of computing gradients over the entire dataset, SGD relies on randomly selected mini-batches to estimate the gradient and update the model parameters to minimize the loss function $f(w)$ (Mahboubi et al., 2024).

The system model adopted in this work follows a simplified setup consistent with Choi et al. (2019). In particular, it assumes an adapted SGD tailored to our system model with independent and identically distributed (IID) data and single-step local updates as shown in Chapter 3. This IID modeling choice allows us to isolate the effects of the communication layer on learning performance, even though this is not the typical FL data distribution according to McMahan et al. (2017).

After finishing the model training, an important step is the analysis of the results. The parameter error norm, given by

$$\text{error norm} = \|w(t) - w^*\|, \quad (1)$$

measures the distance between the estimated model parameters and the optimal solution, and is widely adopted as a standard convergence indicator in stochastic gradient-based algorithms (Li, 2021). In the context of linear regression with quadratic loss, error norm isolates the convergence behavior of the algorithm from observation noise, making it particularly suitable for controlled simulation studies with IID data, such as the IoT learning scenarios considered in this work.

Fig. 1 illustrates the iterative FL process considered in this work. A base station (BS) broadcasts a global model to all devices within a single cell. Each device then




Figure 1 – The iterative Federated Learning process within a single-cell network. The sequence includes the broadcast of the global model, local training at each device, and the transmission of local updates back to the base station for aggregation.

performs local training using its own data and transmits the updated model parameters back to the BS. Upon receiving the local updates, the BS aggregates them to refine the global model, which is then broadcast again for the next iteration.

Beyond the choice of the learning algorithm, the performance of FL systems critically depends on reliable uplink communication, since the BS can aggregate only the updates that are successfully received. Consequently, the design of the uplink transmission scheme is a key factor for overall system performance, especially in IoT scenarios (Elzemity et al., 2024).

The most mature applications of FL are related to mobile devices, such as next-word prediction and on-device personalization, which motivated the original development of the framework (McMahan et al., 2017). In contrast, FL deployments over IoT networks remain an active area of research. Among the many open research challenges are the application of FL in crop yield and flood prediction (Dey et al., 2025; Mahir et al., 2025), topics that align well with the main focus of the CNPq/INCT STREAM (Signal Processing and TRansmission for Environmental Analysis and Monitoring) project (INCT STREAM, 2025), which partially supported this work.

The stringent constraints of IoT devices in terms of energy, computational capability, and memory require the adoption of lightweight communication strategies. In this context, the design of appropriate MAC (Medium Access Control) protocols becomes essential, as discussed in the following section.

2.2 RANDOM ACCESS FOR FEDERATED LEARNING

The performance of FL systems depends critically on the uplink communication scheme, as local updates must be reliably transmitted from devices to the BS for aggregation. In massive IoT deployments, this is particularly challenging: devices are often resource-constrained, may remain in low-power states for extended periods, and exhibit sporadic transmission patterns (Adam et al., 2024; Chen, H. et al., 2022).

Traditional synchronized protocols such as TDMA and FDMA offer collision-free operation but require strict coordination, which increases complexity and energy consumption (Lun et al., 2024; Shah et al., 2021). In contrast, grant-free RA protocols are better suited for FL over IoT, as they allow devices to transmit without prior authorization, reducing signaling overhead and supporting large numbers of devices with sporadic traffic (Silva, R. V. da et al., 2022; Ding et al., 2022; Tian et al., 2023).

Among RA protocols, ALOHA and its variants have emerged as attractive options for FL applications due to their simplicity and low implementation cost. This section reviews the key ALOHA-based schemes relevant to this work: The SA baseline and CRDSA. Although other ALOHA-based protocols such as Diversity Slotted ALOHA (DSA) and Irregular Repetition Slotted ALOHA (IRSA) could also be investigated, throughput performance and implementation complexity were the main factors guiding the selection of CRDSA.

2.2.1 Slotted ALOHA

In SA, time is divided into discrete slots, and devices are constrained to begin transmissions only at slot boundaries. This synchronization reduces the vulnerable period compared to pure ALOHA, where packets can arrive at any time and collide if their transmission intervals overlap (Kleinrock et al., 1975).

Let G denote the normalized offered load, defined as the average number of transmission attempts per slot. Under the assumption that transmission attempts follow a Poisson process, the probability of successful transmission (i.e., no collision) is e^{-G} , and the throughput T , defined as the average number of successful transmissions per slot, is given by (Bellouch et al., 2022):

$$T(G) = Ge^{-G}. \quad (2)$$

The throughput is maximized at $G = 1$, yielding $T_{\max} = 1/e \approx 0.368$, which corresponds to approximately 37% channel utilization (Seo et al., 2025).

In the multichannel setting later adopted in this work, the system comprises M orthogonal channels, and the offered load can be expressed as $G = N/M$, where N is the total number of transmission attempts. SA serves as the baseline RA protocol for FL, as studied in Choi et al. (2019).

2.2.2 Contention Resolution Diversity Slotted ALOHA

CRDSA extends SA by combining packet replication with SIC at the receiver (Kumazawa et al., 2023). Each device transmits R replicas of the same packet in distinct, randomly selected slots. Each replica contains pointers indicating the positions of the other replicas within the frame (Gu et al., 2024).

This structure enables iterative collision resolution: when a replica is successfully decoded in a collision-free slot, the receiver uses the embedded pointers to locate and cancel the corresponding replicas from other slots. This cancellation may reveal previously collided packets, which can then be decoded, triggering further cancellations. The process repeats until no additional packets can be recovered.

Figure 2 illustrates the SIC process. Initially, collisions occur in channels 1, 3, and 4. Device 3 transmits two replicas: once its replica in channel 2 is decoded, the replica in channel 4 can be canceled. This reveals device 1's second replica, which, once decoded, allows cancellation of its first replica in channel 1. The process continues until no more packets can be resolved.¹




Figure 2 – Packet transmission and collision resolution in CRDSA. During iteration t , devices send updates to the BS. The SIC process iteratively cancels replicas of decoded packets, enabling further packet recovery.

Due to the correlations introduced by SIC, the throughput of CRDSA cannot be expressed in closed form and is typically characterized through density evolution

¹Although ALOHA-based protocols allow multiple users to select the same channel, thereby causing collisions, the orthogonality of the channel structure remains preserved. A packet transmitted within a given channel does not interfere with transmissions in other channels.

analysis or simulation (Dossa et al., 2024). Nevertheless, it can be approximated as:

$$T(G) \approx G P_{\text{SIC}}(G, R), \quad (3)$$

where $P_{\text{SIC}}(G, R)$ represents the probability that at least one replica is successfully recovered through interference cancellation. For a given number of replicas R , there exists an optimal operating load G_R^* that maximizes throughput. As shown in chapter 3, the optimal values for R and G will be obtained by simulation.

The key advantage of CRDSA for FL is that it increases the fraction of local updates successfully received at the BS per iteration, thereby accelerating model convergence. The mathematical formulation of CRDSA within the FL framework, including the derivation of access probabilities and feedback control mechanisms, is developed in Chapter 3.

2.3 RELATED WORK

The integration of FL with RA protocols is a recent research area, and several studies have explored different aspects of this problem.

2.3.1 Baseline Work

The work by Choi et al. (2019) serves as the primary baseline for this dissertation. The authors compare multichannel SA with a polling-based access scheme for FL systems. While polling avoids collisions by sequentially requesting updates from devices, it wastes resources when devices are not ready to transmit, resulting in slower convergence. In contrast, SA enables opportunistic transmissions, achieving faster convergence on average, though at the expense of limited throughput under high device density. This work provides a simulation framework for studying ALOHA-based communication in FL and demonstrates the trade-offs between collision-free and contention-based access schemes.

A subsequent study by Rafael Valente da Silva et al. (2022) extends the framework of Choi et al. (2019) by considering multiple learning models with different access probabilities. The authors propose an adaptive communication strategy that selects the appropriate access scheme based on system conditions, demonstrating improved transmission efficiency in congested IoT scenarios. Their work also validates the framework using both synthetic IID data and the Modified National Institute of Standards and Technology (MNIST) dataset for image classification.

2.3.2 Alternative Approaches

Several other works have explored different aspects of FL communication. Hongda Wu et al. (2022) addresses accuracy degradation caused by non-IID data dis-

tributions. The authors propose the Federated Probabilistic Node Selection (FedPNS) framework, which dynamically adjusts device selection probabilities based on their contributions to filter out harmful local updates. While this approach achieves higher accuracy than FedAvg, it introduces additional computational overhead.

Rodolfo Viturino Nogueira Da Silva et al. (2025) introduces a hierarchical FL architecture where devices transmit to cluster heads via D2D communication, and cluster heads forward aggregated updates to the BS using multichannel ALOHA. This approach reduces latency and balances communication loads, though it faces challenges with coordination overhead in dynamic channels.

Salama et al. (2023) investigates decentralized FL (DFL), where devices collaborate directly with neighbors without a central server. Using SA for inter-device communication in mesh topologies, DFL achieves similar accuracy to centralized FL while reducing communication costs and eliminating single points of failure.

More recently, Yu Heng Wu et al. (2025) proposes an Age of Gradient (AoG) policy that prioritizes transmissions based on gradient freshness. By combining gradient sparsification with adaptive transmission probabilities, this approach reduces collisions and enhances convergence, though at the cost of increased protocol complexity.

2.3.3 Positioning of This Work

The works reviewed above highlight a common limitation: SA, while simple and effective for moderate device densities, suffers from throughput degradation under heavy traffic loads due to frequent collisions. This motivates the investigation of CRDSA as an enhanced RA protocol for FL.

Building on the framework established by Choi et al. (2019), this dissertation evaluates CRDSA under varying conditions, including the number of available channels and replica configurations. The goal is to assess how much the higher throughput achieved by CRDSA translates into faster FL convergence and lower error norms, while maintaining the lightweight, grant-free nature essential for IoT deployments.

Table 2.1 summarizes the key characteristics of the most related works.

Table 2.1 – Summary of related work on FL with RA protocols.

Reference	Access Protocol	FL Setting	Main Contribution
Choi et al. (2019)	Multichannel SA, Polling	Centralized, IID	Baseline comparison of SA vs. polling for FL
Rafael Valente da Silva et al. (2022)	Multichannel ALOHA	Centralized, IID/non-IID	Adaptive access probability optimization
Hongda Wu et al. (2022)	—	Centralized, non-IID	Probabilistic node selection (FedPNS)
Rodolfo Viturino Nogueira Da Silva et al. (2025)	Multichannel ALOHA	Hierarchical	Clustered FL with D2D communication
Salama et al. (2023)	SA	Decentralized	Mesh-based FL without central server
Yu Heng Wu et al. (2025)	RA with AoG	Centralized	Age-aware transmission policy
This work	CRDSA	Centralized, IID	CRDSA for FL with SIC-based collision resolution

3 SYSTEM MODEL

The FL process follows an iterative transmission cycle. As illustrated in Fig. 1, each iteration t consists of four main phases: global model download, local model update, local model upload, and global model aggregation (Zhang et al., 2025). The system is modeled as a single-cell topology with K devices, indexed as $k \in \{0, 1, \dots, K-1\}$, a single BS, and M orthogonal channels connecting the devices to the BS. In massive IoT scenarios, we typically assume $K \gg M$, necessitating efficient medium access control.

Following Choi et al. (2019) and Valente da Silva et al. (2023), we assume a common ground-truth parameter vector w^* that the system aims to learn. The k -th device holds a local dataset $D_k = \{x_k, y_k\}$, where x_k is the input and the output y_k is generated according to the linear model:

$$y_k = x_k^T w^*. \quad (4)$$

The global objective is to find an estimate w that minimizes the average loss across all devices, according to the adapted SGD:

$$\min_w \frac{1}{K} \sum_{k=1}^K f_k(w). \quad (5)$$

Specifically, we consider a linear regression setting with a quadratic loss function:

$$f_k(w) = \ell(w, x_k, y_k) = \frac{1}{2} (x_k^T w - y_k)^2. \quad (6)$$

In FL, raw datasets are never exchanged; instead, devices compute local updates and transmit them to the BS. Therefore, we assume that in any given iteration, a device successfully computes its local update with probability p_{comp} . This captures the stochastic nature of device availability and computational resources in IoT networks.

For a computationally active device k , the local update at iteration t is obtained as:

$$w_k(t+1) = w(t) - \mu_w \nabla f_k(w(t)) = w(t) - \mu_w (x_k^T w(t) - y_k) x_k, \quad (7)$$

where $\mu_w > 0$ is the learning rate and the gradient is computed using a single training sample (x_k, y_k) .

After the local updates are computed, devices attempt to transmit their models over the M shared channels. The BS aggregates only the updates that are successfully received (i.e., those that do not suffer from unresolvable collisions). Let $\mathcal{S}_t \subseteq \{1, \dots, K\}$ be the set of devices whose updates are successfully decoded by the BS at iteration t . The BS refines the global model according to:

$$w(t+1) = \frac{1}{|\mathcal{S}_t|} \sum_{k \in \mathcal{S}_t} w_k(t+1). \quad (8)$$

If no updates are received ($|S_t| = 0$), the global model remains unchanged, such that $w_{t+1} = w(t)$.

3.1 FEDERATED LEARNING OVER POLLING

In the Polling approach, the BS acts as a central coordinator that eliminates collisions by sequentially scheduling device transmissions. As illustrated in Fig. 3, the BS manages access to the M orthogonal channels by specifically requesting updates from a subset of devices. While this deterministic approach effectively prevents interference, its performance is highly sensitive to the stochastic availability of devices.




Figure 3 – Polling Protocol: The BS polls devices sequentially. In this framework, the uplink transmission itself serves as an implicit acknowledgment; if a device is computationally active and polled, it transmits its update. If it is inactive, the assigned channel remains idle.

In each iteration t , the BS selects a subset of devices to be polled, where the size of the set is limited by the available channels M . The set of devices that successfully contribute to the global model is restricted to those that were both selected by the BS and were computationally active during that iteration. Consequently, the expected number of successfully aggregated updates in a single polling round is $p_{\text{comp}}M$, which corresponds to the product of the number of channels M and the probability p_{comp} that a device successfully computes its local update.

Although Polling effectively eliminates collisions, it may result in slower convergence compared to random access schemes. This performance degradation stems from the following problem: if a polled device has not completed its computation (with probability $1 - p_{\text{comp}}$), the channel resource remains unused. In contrast, contention-based schemes allow other ready devices to attempt transmission in those resources. Thus, in IoT scenarios where p_{comp} is low or K is very large, the rigid sequential nature of Polling becomes a bottleneck for the learning process.

3.2 FEDERATED LEARNING OVER SLOTTED ALOHA

In contrast to the centralized nature of Polling, the decentralized framework of SA allows multiple devices to attempt transmission independently. This contention-based approach can lead to a higher average number of successfully aggregated updates per round by utilizing slots that would otherwise remain idle in a polling schedule. Following the framework in Choi et al. (2019), we model device availability through two distinct layers. First, a device must be computationally active to produce a fresh update, a state modeled by the probability p_{comp} . Second, even with a valid update, the device must contend for channel access based on a dynamic access probability p_k , which is adapted according to the relevance of its local update and feedback from the BS.

Consequently, for a device k to successfully contribute to the global model at iteration t , three events must occur: (i) the device must be computationally available (with probability p_{comp}); (ii) the device must decide to transmit (with probability p_k); (iii) The transmission must not suffer an unresolvable collision. Because of these combined constraints, the set of successfully received updates $\mathcal{S}_t$ is typically a subset of the K devices.

To determine the optimal p_k , we first define the probability q_k that the BS successfully receives a local update from user k . In an SA system with M channels, this probability depends on both the access decision of device k and the potential collisions with other devices. Specifically, device k 's transmission succeeds if no other device transmits in the same channel. The probability that no other device transmits on a given channel is approximated as (Choi et al., 2019):

$$q_k = p_k \prod_{n \neq k} \left(1 - \frac{p_n}{M}\right) \leq p_k e^{-\frac{\sum_{n \neq k} p_n}{M}} \leq p_k e^{-\frac{P - p_k}{M}}, \quad (9)$$

where $P = \sum_{k=1}^K p_k$ denotes the aggregate access probability (or total offered load). For sufficiently large K and small values of p_n/M , the bounds are tight, and we can approximate:

$$q_k \approx p_k e^{-\frac{P}{M}}. \quad (10)$$

To understand the overall system performance, we aggregate the success probabilities across all users. Let $Q = \sum_{k=1}^K q_k$ denote the expected number of successfully received updates per iteration. Using (10), we obtain:

$$Q = \sum_{k=1}^K p_k e^{-\frac{P}{M}} = e^{-\frac{P}{M}} \sum_{k=1}^K p_k = P e^{-\frac{P}{M}}. \quad (11)$$

The function $f(P) = P e^{-P/M}$ is concave in P and achieves its maximum when $\frac{\partial f}{\partial P} = e^{-P/M} - \frac{P}{M} e^{-P/M} = 0$, which yields $P = M$. Substituting this optimal load into the expression above:

$$Q_{\max} = M e^{-1} \approx 0.368M. \quad (12)$$

Therefore, the expected number of successfully received updates per iteration is bounded by:

$$Q = \sum_{k=1}^K q_k \leq Me^{-1}. \quad (13)$$

This fundamental result shows that the system can aggregate at most Me^{-1} updates per iteration. To achieve this theoretical maximum, the aggregate access probability must satisfy the capacity constraint:

$$\sum_{k=1}^K p_k = M. \quad (14)$$

When this condition is satisfied, the success probability in (10) simplifies to:

$$q_k = p_k e^{-1}, \quad (15)$$

which is a key relationship that will be used in the optimization problem formulated next.

Following Choi et al. (2019), we now formulate an optimization problem that leverages the success probabilities to improve the aggregation quality. The motivation for this optimization stems from a fundamental observation: not all local updates are equally informative. Some devices may compute gradients with large magnitudes, indicating significant information content, while others produce negligible updates. To improve learning efficiency, we should prioritize transmissions from devices with more informative updates.

To quantify the impact of unsuccessful transmissions on the aggregation process, consider the aggregation error. Let $a = \sum_{k=1}^K w_k$ denote the unnormalized aggregation of all local updates (as if all devices transmitted successfully), and let $u = \sum_{k=1}^K w_k \delta_k$ denote the actual aggregated update where $\delta_k \in \{0, 1\}$ indicates whether the BS successfully received the update from device k . The error incurred when device k 's update is not received is proportional to the norm of that update, $\|w_k\|$. Thus, we can bound the aggregation error as:

$$\mathbb{E}[\|a - u\| \mid W] = \mathbb{E}\left[\left\|\sum_{k=1}^K w_k(1 - \delta_k)\right\| \mid W\right] \leq \sum_{k=1}^K a_k \mathbb{E}[1 - \delta_k \mid w_k], \quad (16)$$

where $a_k = \|w_k\|$ represents the importance or relevance of device k 's update. Using the inequality $1 - x \leq e^{-x}$ for $x \in (0, 1)$, we obtain:

$$\mathbb{E}[\|a - u\| \mid W] \leq \sum_{k=1}^K a_k e^{-q_k}, \quad (17)$$

where $q_k = \mathbb{E}[\delta_k \mid w_k]$ is the probability that device k 's transmission is successful.

To minimize the upper bound on the aggregation error, we formulate the following optimization problem. Following the constraint from equation (14), we know that the aggregate success probabilities are limited by $\sum_{k=1}^K q_k = Me^{-1}$. Therefore, the objective is to allocate this limited success probability across devices in a way that prioritizes updates with high information content:

$$\min_{\{q_k\}} \sum_{k=1}^K a_k e^{-q_k}, \quad (18)$$

subject to the aggregate capacity constraint $\sum_{k=1}^K q_k = Me^{-1}$ and the box constraints $q_k \in [0, e^{-1}]$. The box constraint arises because from (10), we have $q_k = p_k e^{-1}$, and since access probabilities are bounded by $p_k \in [0, 1]$, we have $q_k \in [0, e^{-1}]$.

If a device is computationally inactive at iteration t (i.e., it cannot compute a local update), we set $a_k = 0$, which effectively prevents it from contending for access and keeps it from influencing the optimization.

The objective function in equation (18) is convex, and the constraints define a convex feasible set, ensuring a unique global optimum. This structure allows us to use standard convex optimization techniques to find the optimal allocation of success probabilities across all devices. Using the method of Lagrange multipliers, we form the Lagrangian:

$$\mathcal{L}(q_1, \dots, q_K, \lambda) = \sum_{k=1}^K a_k e^{-q_k} + \lambda \left(\sum_{k=1}^K q_k - Me^{-1} \right), \quad (19)$$

where λ represents the sensitivity of the aggregation error to the channel capacity constraint (Boyd et al., 2004). Applying the Karush-Kuhn-Tucker (KKT) conditions, we take the partial derivative with respect to q_k :

$$\frac{\partial \mathcal{L}}{\partial q_k} = -a_k e^{-q_k} + \lambda = 0 \implies q_k = \ln a_k - \ln \lambda. \quad (20)$$

Projecting this onto the feasible interval $[0, e^{-1}]$ and returning to p_k via (10), we obtain the optimal access probability:

$$p_k^{*(\text{SA})} = \left[e \ln(a_k) - \psi \right]_0^1, \quad (21)$$

where $\psi = e \ln \lambda$ and $[\cdot]_0^1$ denotes truncation to the interval $[0, 1]$. This result demonstrates that updates with higher norms (a_k) are prioritized for transmission, while weaker updates are suppressed to reduce contention.

To maintain the system at the optimal operating point, the BS must dynamically calculate and feedback parameter ψ . Let $\hat{P}_t$ be the estimated aggregate access probability at iteration t , calculated as:

$$\hat{P}_t = \sum_{k=1}^K s_{k,t}, \quad s_{k,t} \in \{0, 1\}, \quad (22)$$

where $s_{k,t} = 1$ if device k attempts a transmission and 0 otherwise. The BS updates the feedback signal using a stochastic gradient descent approach:

$$\psi_{t+1}^{(\text{SA})} = \psi_t^{(\text{SA})} + \mu_\psi (\hat{P}_t - M), \quad (23)$$

where μ_ψ is the step size. In this multichannel model, the target load is the total number of channels M , which corresponds to a per-channel offered load of $G = 1$. After each iteration, the BS broadcasts the updated global model $w(t+1)$ and feedback ψ_{t+1} to all devices, closing the control loop.

3.3 FEDERATED LEARNING OVER CRDSA

CRDSA enhances classical SA by incorporating time diversity and iterative processing. Its operation is governed by two fundamental principles: (i) each device transmits R replicas of its packet in randomly selected slots (channels) within a frame, and (ii) the BS utilizes iterative SIC to resolve collisions. Once a single replica is successfully decoded, its interference contribution is reconstructed and subtracted from the other slots where the remaining replicas were transmitted (Kumazawa et al., 2023). This process, illustrated in Fig. 2, significantly improves the packet loss rate and enhances overall system capacity compared to traditional random access schemes.

In a practical CRDSA deployment, the BS manages medium access by storing the entire received signal of a temporal frame within a digital buffer. Each device independently transmits R replicas of its local update; each packet header contains specific metadata, referred to as *pointers*, which indicate the precise temporal positions of its twin replicas within the frame. This signaling mechanism enables the BS to resolve collisions through a physical-layer iterative process.

The core of the decoding logic lies in the iterative SIC procedure. When a replica is successfully decoded, the BS extracts the pointer information to identify the locations of the user's other replicas. The receiver then reconstructs the modulated signal of the decoded user and subtracts it from the buffered signal at the corresponding slot positions. This unveils previously collided packets, potentially rendering them resolvable in subsequent iterations Gu et al. (2024). To evaluate this system, we represent this physical buffering and subtraction logic through the structured two-dimensional transmission matrix and algorithmic flow detailed in Algorithm 1.

Algorithm 1 CRDSA Transmission Matrix Construction and SIC Decoding

Require: Number of users K , number of slots M , number of replicas per user R , slot-allocation matrix $C \in \{1, \dots, M\}^{K \times R}$

Ensure: Set $\mathcal{S}$ of successfully decoded users

```

1: Initialize transmission matrix  $A \leftarrow 0_{K \times M}$  ▷ CRDSA matrix
2: for  $n = 1$  to  $K$  do
3:   for  $r = 1$  to  $R$  do
4:      $s \leftarrow C_{n,r}$  ▷ slot index selected for this replica
5:      $A_{n,s} \leftarrow 1$  ▷ user  $n$  transmits in slot  $s$ 
6:   end for
7: end for
8:  $\mathcal{S} \leftarrow \emptyset$  ▷ set of users decoded successfully
9: while true do
10:   $\mathcal{V} \leftarrow \{m \in \{1, \dots, M\} : \sum_{n=1}^K A_{n,m} = 1\}$  ▷ slots with exactly one packet
11:  if  $\mathcal{V} = \emptyset$  then
12:    break ▷ no more resolvable packets
13:  end if
14:   $\mathcal{U} \leftarrow \emptyset$  ▷ temporary set of valid users
15:  for each  $m \in \mathcal{V}$  do
16:    find  $u$  such that  $A_{u,m} = 1$  ▷ user transmitting in slot  $m$ 
17:     $\mathcal{U} \leftarrow \mathcal{U} \cup \{u\}$ 
18:  end for
19:   $u^* \leftarrow \min(\mathcal{U})$  ▷ select user with minimum index
20:   $\mathcal{S} \leftarrow \mathcal{S} \cup \{u^*\}$  ▷ store successful user
21:   $A_{u^*,m} \leftarrow 0$  for all  $m = 1, \dots, M$  ▷ remove all replicas of user  $u^*$ 
22: end while

```

As established in Chapter 2, the offered load is defined as $G = N/M$. While SA is limited by a maximum throughput at $G = 1$, CRDSA shifts the optimal operating point towards higher loads by effectively resolving collisions that would otherwise result in packet loss. As shown in Fig. 4, increasing the replication factor to $R = 3$ allows the system to achieve superior throughput, providing a more robust communication layer for FL updates.

To integrate this into the FL framework, we adapt the feedback-based access control to target the specific optimal operating point of CRDSA, denoted as G_R^* . By modifying the feedback update rule:

$$\psi_{t+1}^{(\text{CRDSA}, R)} = \psi_t^{(\text{CRDSA}, R)} + \mu_\psi (\hat{P}_t - G_R^* M) \quad (24)$$

we ensure the system maintains an aggregate load that maximizes update success rates. The resulting access probability:

$$p_k^{*(\text{CRDSA}, R)} = \left[e^{G_R^* \ln(a_k)} - \psi^{(\text{CRDSA}, R)} \right]_0^1 \quad (25)$$




Figure 4 – Throughput versus normalized offered load for Slotted ALOHA and CRDSA with different replication factors $R \in \{2, 3, 4\}$.

optimally balances the channel load across the M orthogonal resources, ensuring that the high-capacity benefits of CRDSA are fully exploited to accelerate the convergence of the global model.

4 RESULTS

To assess the learning performance of the proposed scheme, we track the evolution of the model parameters during training using the error norm metric, $\|w(t) - w^*\|$, which measures the distance between the current iteration and the optimal solution. The channel noise is neglected, i.e., the SNR is assumed to approach infinity, and successive interference cancellation (SIC) is considered ideal and lossless. Moreover, unless otherwise stated, the system parameters are set to $K=300$ devices and $M=40$ orthogonal channels, with learning rates $(\mu_w, \mu_\psi) = (0.01, 0.1)$, $p_{\text{comp}} = 0.1$, and for the CRDSA scheme, each active device transmits $R=3$ replicas of its local update, which according to our previous demonstration in chapter 3, is the number of replicas that achieves the best result despite the other parameters, with $G_R^* = 0.67$ ¹. The ground truth w^* and the data input x_k at each device are modeled as independent Gaussian random vectors of dimension $L \times 1$, with $w^* \sim \mathcal{N}(0, I)$ and $x_k \sim \mathcal{N}(0, I)$, where I denotes the $L \times L$ identity matrix. In all experiments, the data length is set to $L = 10$. The proposed strategy is compared against the Polling and SA baselines introduced in (Choi et al., 2019). The SA system applies the access probabilities given in (21), whereas the CRDSA scheme adopts the access probabilities defined in (25).

4.1 NUMERICAL RESULTS

An essential factor in the performance is the aggregate access probability $\hat{P}_t$, which reflects the fraction of devices transmitting at each iteration. Fig. 5 shows convergence for both CRDSA, following (25), and SA according to (21). With CRDSA, the optimal load parameter G_R^* drives the system toward the desired operating point, enhancing learning efficiency. For $R=3$ replicas, the CRDSA convergence stabilizes with a $\hat{P}_t$ at about 67% lower than the optimal SA load, as desired according to Fig. 4.

¹The adaptive feedback mechanism operates as a load control strategy, regulating the offered load G to maintain the system always close to the throughput-optimal operating point of CRDSA even under practical conditions. Specifically, the access probability is dynamically adjusted so that the effective traffic load matches the value that maximizes throughput according to the analytical model, where the number of channels M is parameterized to represent the optimal CRDSA load.




Figure 5 – Convergence of the access probability as a function of the number of iterations for CRDSA and SA. Raw results and their moving average are shown, where the moving average of the true aggregate access probability $\hat{P}_t$ is computed over a sliding window of the most recent t iterations, such that at each iteration the newest value of $\hat{P}_t$ is included and the oldest is discarded.

Fig. 6 illustrates the error norm under different replica configurations for CRDSA, while comparing to that of the baselines. The best performance is obtained when $R = 3$, consistent with Section 3.3. For this configuration, FL over CRDSA achieves considerable performance improvements over SA and Polling. It can also be observed that from $R = 1$ to $R = 3$ there is a decreasing trend in the error norm, indicating a progressive improvement in convergence as additional replicas are introduced. However, when increasing the number of replicas beyond $R = 3$, the performance starts to degrade. This occurs because an excessive number of replicas increases channel occupancy and contention, reducing the effectiveness of interference cancellation and, consequently, impairing the convergence of the learning process.




Figure 6 – Impact of the number of replicas on the error norm $\|w(t) - w^*\|$ at iteration 100.

With the optimal replica configuration, Fig. 7 presents the performance metric Error Norm versus the numbers of channels M . When $M < 8$, CRDSA yields inferior performance compared to SA, as the limited number of channels prevents effective SIC. As M increases, SIC becomes more efficient, and the performance of CRDSA improves accordingly. The proposed scheme surpasses both SA and Polling, with performance gains that increase consistently with the number of channels, reaching their maximum at $M = 40$. These results indicate that the performance gains achieved by CRDSA increase with the number of available channels, as the effectiveness of successive interference cancellation improves and a larger fraction of local updates can be successfully aggregated at each round.




Figure 7 – Error norm $\|w(t) - w^*\|$ at iteration 100 as a function of the number of available orthogonal channels, for CRDSA, Slotted ALOHA, and Polling.

To evaluate model learning performance, after set a fair parameter value for the number of channels M , Fig. 8 depicts the evolution of the error norm over 500 iterations. For instance, to reach an error norm of 10^{-5} , the proposed CRDSA-based scheme converges significantly faster than the baseline methods, requiring approximately 50% fewer iterations than SA and around 85% fewer iterations than Polling. This accelerated convergence can be attributed to the higher fraction of successfully aggregated updates per iteration enabled by CRDSA. By mitigating packet collisions through successive interference cancellation, the proposed scheme preserves a larger number of informative local updates at each round, effectively reducing the variance of the aggregated gradient. As a result, the global model progresses more steadily toward the optimal solution, leading to a faster decay of the error norm when compared to Slotted ALOHA and Polling, whose convergence is more strongly affected by collision losses and sequential access constraints, respectively.




Figure 8 – Convergence measured by the error norm $\|w(t) - w^*\|$ versus the number of iterations, comparing CRDSA, Slotted ALOHA, and Polling.

Although the results in Fig. 8 indicate that the proposed scheme achieves superior learning performance, massive IoT deployments typically involve a very large number of connected devices. Therefore, Fig. 9 compares the performance of Polling, SA, and CRDSA under increasing user density. The results show that CRDSA attains improved performance for scenarios with more than $K = 1000$ users. However, as the number of users increases, the level of contention also rises, which adversely impacts the learning process. As previously demonstrated in this chapter, the CRDSA-based approach achieves a lower error norm only when a sufficient number of channels (resources) is available. Consequently, accommodating an increasing number of users or other degrading system parameters requires either a proportional increase in the number of available channels or results in a significant degradation of learning performance. This behavior highlights the trade-off between user density and available communication resources in random access-based federated learning systems. For moderate numbers of users, the interference cancellation capabilities of CRDSA allow the system to sustain a high fraction of successfully aggregated updates, resulting in a lower error norm compared to Slotted ALOHA and Polling. However, when the number of users becomes excessively large relative to the number of channels, the collision rate increases beyond the resolution capability of successive interference cancellation, leading to a gradual degradation in convergence performance. These results indicate that

the gains provided by CRDSA are strongly coupled to the ratio between user density and channel availability.




Figure 9 – Error norm $\|w(t) - w^*\|$ at iteration 100 as a function of the number of users, for CRDSA, Slotted ALOHA, and Polling.

Finally, the impact of variations in the probability of computation is evaluated in Fig. 10, where p_{comp} ranges from 0 to 1. The proposed CRDSA-based scheme outperforms SA for $p_{\text{comp}} < 0.3$ and outperforms Polling for $p_{\text{comp}} < 0.43$. Recall that p_{comp} denotes the probability that a device successfully computes its local update. Since a device can transmit only after completing the computation, the access probability must satisfy $p \leq p_{\text{comp}}$ and is applied independently across all K devices. As noted in Choi et al. (2019), the performance of adaptive access schemes like SA and CRDSA appears to degrade as p_{comp} increases. This occurs because a higher p_{comp} means more devices are ready to transmit simultaneously, which intensifies contention and requires more iterations for the feedback parameter ψ in (23) or (24) to converge to a stable value that effectively manages network load. In network scenarios where most devices successfully complete their updates at each iteration, p_{comp} may approach high values. However, this is generally not the case in typical IoT environments.




Figure 10 – Error norm $\|w(t) - w^*\|$ at iteration 100 versus the probability p_{comp} at iteration 100.

5 CONCLUSIONS AND FUTURE WORK

This work evaluated the integration of the CRDSA protocol into Federated Learning systems operating in massive IoT scenarios, with particular emphasis on communication efficiency under high device density. The main objective of evaluating the suitability of CRDSA for FL-enabled IoT networks was achieved, as the results consistently demonstrated superior convergence behavior and lower error norms when compared to SA and Polling-based baselines, provided that sufficient communication resources are available to enable effective interference cancellation. The analysis confirmed that CRDSA constitutes a scalable alternative for uplink communication in decentralized learning systems with higher throughput when compared to SA.

We have obtained customized CRDSA equations tailored to the federated learning framework adopted in this work. Starting from the classical random access formulation, this dissertation derived access probability and feedback control expressions compatible with the FL aggregation process, explicitly incorporating learning-related parameters such as the impact of local updates and the optimal operating load G_R^* . The resulting mathematical model is a great step towards random access theory and learning dynamics bridge, allowing CRDSA to be embedded into the FL system model.

The communication performance of CRDSA under variations of key system parameters, including the number of orthogonal channels, the number of training iterations, and the local computation probability p_{comp} , was also depicted. The analysis revealed that CRDSA benefits significantly from an increased number of channels, which enhances the effectiveness of successive interference cancellation and leads to faster convergence rates. Additionally, the results showed that CRDSA is particularly robust in low-computation-availability regimes, where packet losses have a more pronounced impact on learning performance. These findings highlight the strong coupling between communication reliability and learning efficiency in FL systems and confirm the importance of adaptive access control mechanisms.

As the main advantages of CRDSA over SA within a federated learning context, the results clearly indicate that CRDSA offers substantial gains by preserving a higher fraction of useful local updates through diversity and successive interference cancellation. Compared to SA, CRDSA achieves lower error norms, faster convergence, and improved resilience to collisions, especially in dense network scenarios, if enough resources are available. In contrast to polling, CRDSA avoids the limitations of sequential access while maintaining low signaling overhead, as only short replica pointers are included in the packet header without increasing the payload size. Furthermore, successive interference cancellation (SIC) is performed at the BS, and its computational complexity does not impact the devices, making CRDSA well suited for large-scale IoT deployments. These improvements demonstrate that CRDSA effectively addresses

key bottlenecks of conventional random access schemes when applied to distributed learning systems.

As discussed in Chapter 2, the analysis in this work was conducted under the assumption of IID data distributions in order to isolate the effects of the medium access layer on learning performance. Nevertheless, the proposed framework and insights are not inherently restricted to IID settings. Although non-IID data distributions are easier to find in IoT environments, and typically introduce higher gradient variance and slower convergence, the communication advantages provided by CRDSA are expected to remain relevant, suggesting that the proposed approach can be extended to more realistic federated learning scenarios.

Finally, future research directions include the investigation of higher-throughput random access protocols, such as Irregular Repetition Slotted ALOHA (IRSA), as well as the extension of the proposed framework to non-IID data distributions and heterogeneous device capabilities. Such studies may further enhance the scalability and robustness of federated learning systems operating over random access channels.

REFERENCES

- ADAM, Mumin; BAROUDI, Uthman. Federated Learning for IoT: Applications, Trends, Taxonomy, Challenges, Current Solutions, and Future Directions. **IEEE Open Journal of the Communications Society**, v. 5, p. 7842–7877, 2024.
- BELLOUCH, Abdessamad; ZAALLOUL, Abdellah; HAQIQ, Abdelkrim. Random Access Mechanism Enhancement Based on a Hybrid ALOHA Protocol Using an Analytical Model. **Journal of ICT Standardization**, River Publishers, v. 10, n. 3, p. 383–410, 2022.
- BOYD, S.; VANDENBERGHE, L. **Convex Optimization**. [S.l.]: Cambridge University Press, 2004. ISBN 9781107394001.
- CELANDRONI, Nedo; SECCHI, Raffaello. Suitability of DAMA and Contention-Based Satellite Access Schemes for TCP Traffic in Mobile DVB-RCS. **IEEE Transactions on Vehicular Technology**, v. 58, n. 4, p. 1836–1845, 2009.
- CHEN, Hao; HUANG, Shaocheng; ZHANG, Deyou; XIAO, Ming; SKOGLUND, Mikael; POOR, H. Vincent. Federated Learning Over Wireless IoT Networks With Optimized Communication and Resources. **IEEE Internet of Things Journal**, v. 9, n. 17, p. 16592–16605, 2022.
- CHEN, Mingzhe; YANG, Zhaohui; SAAD, Walid; YIN, Changchuan; POOR, H. Vincent; CUI, Shuguang. A Joint Learning and Communications Framework for Federated Learning Over Wireless Networks. **IEEE Transactions on Wireless Communications**, v. 20, n. 1, p. 269–283, 2021.
- CHOI, Jinho; POKHREL, Shiva. Federated Learning with Multichannel ALOHA. **IEEE Wireless Communications Letters**, dez. 2019.
- DEY, Tanushree; BERA, Somnath; MUKHERJEE, Anwesha; DE, Debashis; BUYYA, Rajkumar. FLYer: Federated Learning-Based Crop Yield Prediction for Agriculture 5.0. **IEEE Transactions on Artificial Intelligence**, v. 6, n. 7, p. 1943–1952, 2025.
- DING, Zhiguo; LE NGO, Phuong; CHOI, Junil. Grant-Free Random Access in Machine-Type Communication: Approaches and Challenges. **IEEE Wireless Communications**, v. 29, n. 1, p. 151–158, 2022.
- DOSSA, Amavi; AMHOUD, El Mehdi. A Duty-Cycle-Efficient Synchronization Protocol for Slotted-Aloha in LoRaWAN. In: IEEE Global Communications Conference (GLOBECOM). [S.l.: s.n.], 2024. p. 1839–1844.
- ELZEMITY, Adel; ARIEF, Budi. Privacy Threats and Countermeasures in Federated Learning for Internet of Things: A Systematic Review. In: IEEE International

Conferences on Internet of Things and IEEE Green Computing & Communications and IEEE Cyber, Physical & Social Computing and IEEE Smart Data and IEEE Congress on Cybermatics. [S.l.: s.n.], 2024. p. 331–338.

GAO, Yansong; KIM, Minki; THAPA, Chandra; ABUADBBA, Alsharif; ZHANG, Zhi; CAMTEPE, Seyit; KIM, Hyounghshick; NEPAL, Surya. Evaluation and Optimization of Distributed Machine Learning Techniques for Internet of Things. **IEEE Transactions on Computers**, v. 71, n. 10, p. 2538–2552, 2022.

GU, Yiyang; XU, Yunlai; ZHANG, Bo; WANG, Ye; YANG, Zhihua. Toward the Random Multiaccess in SloT: A Generalized-Deduplication-Based CRDSA Mechanism. **IEEE Internet of Things Journal**, v. 11, n. 11, p. 20207–20222, 2024.

HUANG, Yuhao; JIAO, Jian; WU, Shaohua; LU, Rongxing; ZHANG, Qinyu. Age-Critical Frameless ALOHA Protocol for Grant-Free Massive Access. In: IEEE Global Communications Conference (GLOBECOM). [S.l.: s.n.], 2021. p. 1–6.

IIDA, Renato F.; TRINDADE, Pedro H. A.; FARIA, Bruno; AGUAYO, Leonardo; WYGLINSKI, Alexander M. SatSysSim: A Novel Event-Driven Simulation Framework for DVB/RCS2 Performance Characterization. **IEEE Access**, v. 10, p. 308–318, 2022.

INCT STREAM. **INCT STREAM – National Institute of Science and Technology**. [S.l.: s.n.], 2025. <https://www.inctstream.org/>. Accessed: Feb. 2025.

KLEINROCK, Leonard; TOBAGI, Fouad A. Packet switching in radio channels: Part I — Carrier sense multiple-access modes and their throughput-delay characteristics. **IEEE Transactions on Communications**, COM-23, n. 12, p. 1400–1416, 1975.

KUMAZAWA, Kansuke; OKADA, Hiraku; MATSUI, Munehiro; TATEGAMI, Koyo; GOTO, Daisuke; ITOKAWA, Kiyohiko; YAMASHITA, Fumihiko; KATAYAMA, Masaaki. Transmission Control Considering Different Available Channels for CRDSA/IRSA in Satellite IoT Systems. **IEEE Access**, v. 11, p. 135377–135384, 2023.

LI, Xiang. Novel Convergence Results of Adaptive Stochastic Gradient Descents. **IEEE Transactions on Image Processing**, v. 30, p. 1044–1056, 2021.

LUN, Yuriy Zacchia; RINALDI, Claudia; D'INNOCENZO, Alessandro; SANTUCCI, Fortunato. Co-Designing Wireless Networked Control Systems on IEEE 802.15.4-Based Links Under Wi-Fi Interference. **IEEE Access**, v. 12, p. 71157–71183, 2024.

MAHBOUBI, Shahrzad; NINOMIYA, Hiroshi. Study on Stochastic Gradient Descent Without Explicit Error Backpropagation with Momentum. In: 2024 IEEE Conference on Artificial Intelligence (CAI). Los Alamitos, CA, USA: IEEE Computer Society, jun. 2024. p. 1357–1358.

- MAHIR, Shahariar Hossain; TASHRIF, Md Tanjum An; KARIM, Md Ahsan; KUNDU, Dipanjali; RAHMAN, Anichur; HAMZA, Md. Amir; AL FARID, Fahmid; MIAH, Abu Saleh Musa; MANSOR, Sarina. Advanced Hydro-Informatic Modeling Through Feedforward Neural Network, Federated Learning, and Explainable AI for Enhancing Flood Prediction. **IEEE Open Journal of the Computer Society**, v. 6, p. 726–738, 2025.
- MCMAHAN, H. B.; MOORE, Eider; RAMAGE, Daniel; HAMPSON, Seth; ARCAS, B.A.Y. Communication-Efficient Learning of Deep Networks from Decentralized Data. In: INTERNATIONAL Conference on Artificial Intelligence and Statistics. [S.l.: s.n.], 2017.
- MIKHAYLOV, Konstantin; STUSEK, Martin; MASEK, Pavel; FUJDIK, Radek; MOZNY, Radek; ANDREEV, Sergey; HOSEK, Jiri. Communication Performance of a Real-Life Wide-Area Low-Power Network Based on Sigfox Technology. In: IEEE International Conference on Communications (ICC). [S.l.: s.n.], 2020. p. 1–6.
- NGUYEN, Dinh C.; DING, Ming; PATHIRANA, Pubudu N.; SENEVIRATNE, Aruna; LI, Jun; VINCENT POOR, H. Federated Learning for Internet of Things: A Comprehensive Survey. **IEEE Communications Surveys & Tutorials**, v. 23, n. 3, p. 1622–1658, 2021.
- PERAZZONE, Jake; WANG, Shiqiang; JI, Mingyue; CHAN, Kevin S. Communication-Efficient Device Scheduling for Federated Learning Using Stochastic Optimization. In: IEEE INFOCOM 2022 - IEEE Conference on Computer Communications. [S.l.: s.n.], 2022. p. 1449–1458.
- SALAMA, Abdelaziz; STERGIOULIS, Achilleas; HAYAJNEH, Ali M.; ZAIDI, Syed Ali Raza; MCLERNON, Des; ROBERTSON, Ian. Decentralized Federated Learning Over Slotted ALOHA Wireless Mesh Networking. **IEEE Access**, v. 11, p. 18326–18342, 2023.
- SEO, Jun-Bae; HU, Yangqian; JIN, Hu; DE, Swades. ALOHA With SIC-Aided Collision Resolution. **IEEE Internet of Things Journal**, v. 12, n. 8, p. 10194–10209, 2025.
- SHAH, A. F. M. Shahan; QASIM, Ahmed Nidham; KARABULUT, Muhammet Ali; ILHAN, Haci; ISLAM, Md. Baharul. Survey and Performance Evaluation of Multiple Access Schemes for Next-Generation Wireless Communication Systems. **IEEE Access**, v. 9, p. 113428–113442, 2021.
- SILVA, Rafael Valente da; CHOI, Jinho; PARK, Jihong; BRANTE, Glauber; SOUZA, Richard Demo. Multichannel ALOHA Optimization for Federated Learning With Multiple Models. **IEEE Wireless Communications Letters**, v. 11, n. 10, p. 2180–2184, 2022.

SILVA, Rodolfo Vitorino Nogueira Da; FLÁVIA DOS REIS, Ana; BRANTE, Glauber; SOUZA, Richard Demo. Hierarchical Federated Learning With Distributed Clustering and Multichannel ALOHA. **IEEE Sensors Journal**, v. 25, n. 7, p. 12128–12142, 2025.

STATISTA. **IoT Connected Devices Worldwide**. [S.l.], 2025. Disponível em: <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/>.

TIAN, Chuan; AMALLADINNE, Vamsi K.; WEI, Yihan; FLETCHER, Alexander K.; CAIRE, Giuseppe. Federated Learning with Unsourced Random Access. **IEEE Transactions on Signal Processing**, v. 71, p. 2904–2917, 2023.

VALENTE DA SILVA, Rafael; LÓPEZ, Onel L. Alcaraz; SOUZA, Richard Demo. Energy-Aware Federated Learning With Distributed User Sampling and Multichannel ALOHA. **IEEE Communications Letters**, v. 27, n. 10, p. 2867–2871, 2023.

WU, Hongda; WANG, Ping. Node Selection Toward Faster Convergence for Federated Learning on Non-IID Data. **IEEE Transactions on Network Science and Engineering**, v. 9, n. 5, p. 3099–3111, 2022.

WU, Yu Heng; ASGARI, Houman; RINI, Stefano; MUNARI, Andrea. Age-of-Gradient Updates for Federated Learning over Random Access Channels. In: IEEE International Conference on Machine Learning for Communication and Networking. [S.l.: s.n.], 2025. p. 1–6.

ZHANG, Hongyi; BOSCH, Jan; OLSSON, Helena Holmström. Enabling efficient and low-effort decentralized Federated Learning with the EdgeFL framework. **Information and Software Technology**, v. 178, p. 107600, 2025. ISSN 0950-5849.