



**UNIVERSIDAD DE MURCIA**  
**ESCUELA INTERNACIONAL DE DOCTORADO**

**TESIS DOCTORAL**

Privacy Preservation Mechanisms in Identity Management with  
Application to IoT

Mecanismos de Preservación de Privacidad en Gestión de  
Identidad con Aplicación a IoT

**D. Jesús García Rodríguez**

**2024**





**UNIVERSIDAD DE MURCIA**  
**ESCUELA INTERNACIONAL DE DOCTORADO**  
**TESIS DOCTORAL**

Privacy Preservation Mechanisms in Identity Management with  
Application to IoT

Mecanismos de Preservación de Privacidad en Gestión de Identidad  
con Aplicación a IoT

Autor: D. Jesús García Rodríguez

Directores: D. Antonio Fernando Skármeta Gómez  
D. Jorge Bernal Bernabé





**DECLARACIÓN DE AUTORÍA Y ORIGINALIDAD  
DE LA TESIS PRESENTADA EN MODALIDAD DE COMPENDIO O ARTÍCULOS PARA  
OBTENER EL TÍTULO DE DOCTOR**

*Aprobado por la Comisión General de Doctorado el 19-10-2022*

D./Dña. Jesús García Rodríguez

doctorando del Programa de Doctorado en

Informática

de la Escuela Internacional de Doctorado de la Universidad Murcia, como autor/a de la tesis presentada para la obtención del título de Doctor y titulada:

Privacy Preservation Mechanisms in Identity Management with Application to IoT / Mecanismos de Preservación de Privacidad en Gestión de Identidad con Aplicación a IoT

y dirigida por,

D./Dña. Antonio Fernando Skarmeta Gómez

D./Dña. Jorge Bernal Bernabé

D./Dña.

**DECLARO QUE:**

La tesis es una obra original que no infringe los derechos de propiedad intelectual ni los derechos de propiedad industrial u otros, de acuerdo con el ordenamiento jurídico vigente, en particular, la Ley de Propiedad Intelectual (R.D. legislativo 1/1996, de 12 de abril, por el que se aprueba el texto refundido de la Ley de Propiedad Intelectual, modificado por la Ley 2/2019, de 1 de marzo, regularizando, aclarando y armonizando las disposiciones legales vigentes sobre la materia), en particular, las disposiciones referidas al derecho de cita, cuando se han utilizado sus resultados o publicaciones.

Además, al haber sido autorizada como compendio de publicaciones o, tal y como prevé el artículo 29.8 del reglamento, cuenta con:

- *La aceptación por escrito de los coautores de las publicaciones de que el doctorando las presente como parte de la tesis.*
- *En su caso, la renuncia por escrito de los coautores no doctores de dichos trabajos a presentarlos como parte de otras tesis doctorales en la Universidad de Murcia o en cualquier otra universidad.*

Del mismo modo, asumo ante la Universidad cualquier responsabilidad que pudiera derivarse de la autoría o falta de originalidad del contenido de la tesis presentada, en caso de plagio, de conformidad con el ordenamiento jurídico vigente.

En Murcia, a 03 de mayo de 2024

Fdo.: Jesús García Rodríguez



# Agradecimientos

Como no puede ser de otra forma, quiero dedicar unas palabras de agradecimiento a todas las personas que me han acompañado y ayudado en el desarrollo de esta tesis doctoral.

En primer lugar, quiero expresar mi más sincero agradecimiento a mis tutores, Antonio y Jorge, por acogerme en el mundo de la investigación, incluso antes de comenzar esta tesis. Gracias por confiar en mí para formar parte de nuestro gran equipo. También por enseñarme a enfrentar las muy diversas facetas necesarias en este mundillo, desde las publicaciones científicas a los proyectos de investigación o innovación. Vuestro apoyo y orientación han sido fundamentales para mi desarrollo académico y personal estos años. Quiero extender también un agradecimiento especial a Stephan, mi tutor durante mi estancia en AIT, por ayudarme a expandir mis horizontes de investigación a un apartado más teórico y sobre todo por el maravilloso trato tanto en ese período como fuera de él.

Gracias a mis colegas en estos años de investigación en la universidad, tanto aquellos con implicación más directa en publicaciones o proyectos, en especial Daniel, Rafa, Agustín o Stefan, como los miembros de Dibulibu, T3 y Pleiades. No se puede pedir un mejor ambiente para completar este arduo camino.

A quien más debo agradecer, por vuestro amor y apoyo incondicional, es a mi familia. Sin vosotros, nada de esto (ni de nada más) habría sido posible. A mis padres, José Luis y Puri, por darme la oportunidad. A mi hermana María José (aunque siempre serás hermana) y a Fran, por engañarme para comenzar este camino, y servirme siempre de ejemplo. También a los que nos dejaron, sin vosotros no sería quien soy hoy, y os llevo siempre en mis recuerdos y mi corazón.

A mis amigos, gracias por vuestra amistad y los momentos compartidos, un contraste de alegría y holgazanería necesario para abordar el trabajo y esfuerzo.

A Raquel, gracias por disfrutar conmigo de los buenos momentos y alegrías y por apoyarme cuando las cosas se torcían tanto en lo personal como lo profesional. El tiempo que hemos compartido ha hecho de estos años de tesis un período más especial y grato de lo que podía imaginar.





# Contents

<b>Resumen</b>	<b>vii</b>
1 Motivación . . . . .	vii
2 Objetivos y metodología . . . . .	ix
3 Resultados . . . . .	x
3.1 Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures . . . . .	xi
3.2 Beyond Selective Disclosure: Extending Distributed p-ABC Imple- mentations by Commit-and-Prove Techniques . . . . .	xii
3.3 To Pass or Not to Pass: Privacy-Preserving Physical Access Control	xiii
3.4 A privacy-preserving attribute-based framework for IoT identity li- fecycle management . . . . .	xiv
4 Conclusiones y trabajo futuro . . . . .	xiv
<b>Abstract</b>	<b>xvii</b>
1 Motivation . . . . .	xvii
2 Objectives and Methodology . . . . .	xix
3 Results . . . . .	xx
3.1 Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures . . . . .	xx
3.2 Beyond Selective Disclosure: Extending Distributed p-ABC Imple- mentations by Commit-and-Prove Techniques . . . . .	xxii
3.3 To Pass or Not to Pass: Privacy-Preserving Physical Access Control	xxii
3.4 A privacy-preserving attribute-based framework for IoT identity lifecycle management . . . . .	xxiii
4 Conclusions and future work . . . . .	xxiv
<b>1 Introduction</b>	<b>1</b>
1.1 Objectives . . . . .	4
1.2 Methodology . . . . .	4
<b>2 Summary of results</b>	<b>6</b>
2.1 Related work . . . . .	6
2.1.1 Digital identity management . . . . .	6
2.1.2 Privacy-preserving attribute-based credentials . . . . .	9
2.1.3 IoT identity management . . . . .	13
2.1.4 Gap analysis . . . . .	15
2.2 Results . . . . .	16
2.2.1 Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures . . . . .	19

2.2.2	Beyond Selective Disclosure: Extending Distributed p-ABC Implementations by Commit-and-Prove Techniques . . . . .	21
2.2.3	To Pass or Not to Pass: Privacy-Preserving Physical Access Control	27
2.2.4	A privacy-preserving attribute-based framework for IoT identity lifecycle management . . . . .	30
<b>3</b>	<b>Conclusions and future works</b>	<b>36</b>
3.1	Future work . . . . .	38
<b>4</b>	<b>Publications composing the doctoral thesis</b>	<b>39</b>
4.1	Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures . . . . .	39
4.2	A privacy-preserving attribute-based framework for IoT identity lifecycle management . . . . .	41
4.3	To pass or not to pass: Privacy-preserving physical access control . . . . .	42
4.4	Beyond Selective Disclosure: Extending Distributed p-ABC Implementations by Commit-and-Prove Techniques . . . . .	43
	<b>Bibliography</b>	<b>63</b>
	<b>Publications</b>	<b>66</b>
	<b>A Abbreviations</b>	<b>67</b>

# List of Figures

- 2.1 Simplified OLYMPUS architecture . . . . . 20
- 2.2 Execution time results of PS-MS with BLS12-461 for issuance and ZK proofs 21
- 2.3 Comparison of execution times between Idemix and PS-MS with BLS12-461 22
- 2.4 Scenario and flows supported by new dp-ABC functionality in contrast to previous solution . . . . . 23
- 2.5 Schematic view of the applications implemented through the commit-and-prove extension . . . . . 24
- 2.6 Execution times for the dp-ABC extended predicate functionalities . . . . . 25
- 2.7 Comparison of the proposed implementation with Idemix . . . . . 25
- 2.8 Biometric-Bound ABC components and communication flow . . . . . 27
- 2.9 Bootstrapping and enrolment in the IoT identity management framework . 31
- 2.10 Operational phase in the IoT identity management framework . . . . . 32
- 2.11 Execution time of wallet methods in Raspberry 4 of our implementation and BBS+ scheme . . . . . 33

# List of Tables

1	Resumen de los resultados de la tesis y su relación con los objetivos conseguidos y publicaciones realizadas . . . . .	xi
2	Summary of results and their link to objectives and publications achieved . . . . .	xxi
2.1	Overhead in token size of the implemented extended predicates . . . . .	26
2.2	Overview of relevant use cases and the most appealing features of the proposed solution . . . . .	26
2.3	Computational overhead of BioABC-ZK instantiation . . . . .	30
2.4	Overview on state-of-art solutions on IoT identity management and which challenges they focus on tackling . . . . .	34
A.1	Abbreviations . . . . .	67



# List of publications

This Doctoral Thesis is presented as a compendium of the following publications, the PhD student being the main author in all of them:

- García-Rodríguez, J., Torres Moreno, R., Bernal Bernabe, J., & Skarmeta, A. (2021). Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures. *Journal of Information Security and Applications*, 62, 102971. <https://doi.org/10.1016/j.jisa.2021.102971>
- García-Rodríguez, J., & Skarmeta, A. (2023). A privacy-preserving attribute-based framework for IoT identity lifecycle management. *Computer Networks*, 236, 110039. <https://doi.org/10.1016/j.comnet.2023.110039>
- García-Rodríguez, J., Krenn, S., & Slamanig, D. (2024). To pass or not to pass: Privacy-preserving physical access control. *Computers & Security*, 136, 103566. <https://doi.org/10.1016/j.cose.2023.103566>
- García-Rodríguez, J., Krenn, S., Bernal Bernabe, J., & Skarmeta, A. (2024). Beyond Selective Disclosure: Extending Distributed p-ABC Implementations by Commit-and-Prove Techniques. *Computers & Security*, 248 , 110498. <https://doi.org/10.1016/j.comnet.2024.110498>

# Resumen

## 1 Motivación

Los servicios digitales se han convertido en un elemento básico de nuestra sociedad. En este entorno es necesario autenticar las diversas entidades digitales y, en particular, a los usuarios. Este requisito no solo abarca la identificación, sino también la autenticación de atributos de identidad como el correo electrónico o la edad de una persona. Así pues, hay que aplicar soluciones de gestión de identidad para lograr un ecosistema seguro. El carácter ubicuo de esta necesidad hace que las propiedades de seguridad y privacidad que consiguen estas soluciones sean más relevantes que nunca.

Los sistemas de identidad tradicionales, muy extendidos por su facilidad de uso, dan lugar a graves problemas de seguridad y privacidad. Por ejemplo, la vulneración de bases de datos de proveedores de servicios [1–3] ha provocado la filtración de miles de millones de registros con información sobre usuarios, como correos electrónicos, contraseñas y datos personales. Los sistemas de inicio de sesión único (*Single Sign-On*, SSO) mitigan el problema reutilizando la información de autenticación a través de proveedores de identidad (IdP) para acceder a múltiples servicios. Sin embargo, esto conlleva problemas de *punto único de fallo* en términos de seguridad y privacidad, culminando en el seguimiento exhaustivo de la actividad online de todos los usuarios. De hecho, estos mecanismos son generalmente ofrecidos por grandes empresas tecnológicas como Google o Facebook, que ya han estado en el punto de mira por faltas contra la privacidad de los usuarios [4].

Estos problemas de seguridad y privacidad no se limitan a las interacciones tradicionales de los usuarios con los servicios digitales. Como destaca el Comité Económico y Social Europeo (CESE), la evolución de las tecnologías digitales y su aplicación a la digitalización de la sociedad brinda enormes oportunidades de crecimiento económico, justicia e inclusión, pero no está exenta de amenazas. Riesgos como violaciones de la privacidad o la introducción de prejuicios y discriminaciones se incrementan considerablemente. Por ello, las políticas y la legislación europeas tienden ahora hacia la soberanía digital<sup>1</sup>. El auge del *Internet de las Cosas* (IoT), donde máquinas o dispositivos son los principales actores en el intercambio de información [5], es uno de los principales exponentes de esta tendencia hacia la digitalización. En 2023, el número de dispositivos IoT conectados alcanzó los 15.000 millones, y las estimaciones previstas para los próximos años reflejan un crecimiento significativo, alcanzando los 30.000 millones de dispositivos en 2030 [6]. Estos

---

<sup>1</sup><https://www.eesc.europa.eu/en/our-work/opinions-information-reports/opinions/digital-identity-data-sovereignty-and-path-towards-just-digital-transition-citizens-living-information-society>

dispositivos tienen características muy variables en cuanto a potencia de cálculo, memoria o limitaciones de consumo. Su ubicuidad y los datos sensibles que se comparten en entornos IoT enfatizan la necesidad de soluciones seguras y con mecanismos de privacidad, particularmente en el apartado de gestión de identidad.

En los últimos años, la tendencia ha virado hacia soluciones más respetuosas con la privacidad, que sitúan a los usuarios (y sus dispositivos) en el centro de los sistemas de gestión de identidad. En este sentido, la emergente auto-soberanía en la gestión de identidad (*Self-Sovereign Identity*, SSI) [7] establece preceptos como la minimización de datos, no trazabilidad, o el consentimiento del usuario. La importancia de estos conceptos se pone de manifiesto en normativas como el RGPD. En particular, la Comisión Europea pretende poner la identidad digital a disposición de todos los ciudadanos de la UE a través del reglamento eIDAS2 [8]. Esta propuesta establece el uso de carteras de identidad (*wallet*) para la autenticación en servicios digitales en toda la UE siguiendo los principios de la auto-soberanía. Por otro lado, el enfoque SSI no solo es apto para individuos, sino que también es aplicable al mundo IoT [9] y, en general, a cualquier servicio u objeto digital.

En la persecución de estos ambiciosos objetivos de privacidad, se ha hecho evidente la necesidad de aplicar herramientas criptográficas específicamente diseñadas, conocidas como tecnologías de mejora de la privacidad (*Privacy-Enhancing Technologies*, PET). Las PET permiten recoger, procesar, analizar o compartir información protegiendo al mismo tiempo la confidencialidad de los datos personales. Existen PET con diversas finalidades como ofuscar datos (mediante manipulación u ocultación), permitir el tratamiento de datos cifrados, o realizar análisis de datos distribuidos. Estas técnicas no pueden considerarse una "panacea" que resuelve todos los problemas de privacidad y protección de datos y su implantación sigue presentando algunas barreras, como su reducida accesibilidad en entornos políticos por su carácter innovador. No obstante, son una herramienta clave para alcanzar el paradigma de la *privacidad desde el diseño*. Además, habilitan nuevas aplicaciones y casos de uso, ya que permiten garantizar los derechos de privacidad al tiempo que se mantiene la utilidad de los datos. Entre las PET existentes, las credenciales basadas en atributos (*privacy-preserving Attribute-Based Credentials*, p-ABC) [10, 11] destacan por su idoneidad para la gestión de identidad de propósito general, proporcionando a los usuarios el control sobre su identidad online. De hecho, se consideran un elemento clave para futuros desarrollos que mejoren la privacidad en la solución europea de identidad digital. Así pues, en esta tesis nos centramos en las p-ABCs y su aplicación para resolver retos clave de la gestión de identidad.

Las p-ABCs permiten generar credenciales digitales que certifican los atributos del usuario. Una vez en posesión del usuario, estas credenciales pueden usarse sin intervención de terceros para derivar *tokens* no vinculables mediante pruebas de conocimiento cero, divulgando selectivamente parte de la información contenida en ellas. En este proceso, el verificador seguirá teniendo garantías formales de que los datos revelados fueron realmente certificados por el emisor. Así pues, las p-ABC pueden aplicarse para realizar la autenticación de atributos de identidad relacionados con procesos de autorización de grano fino manteniendo al mismo tiempo altas garantías de privacidad frente al proveedor del servicio. Además, se mitiga el problema del seguimiento de usuarios por parte de



entidades centrales como el proveedor de identidad.

A pesar de los esfuerzos y avances, el campo de la gestión de identidad aún plantea diversos retos en términos de seguridad y preservación de la privacidad. En primer lugar, la aplicabilidad de los sistemas de gestión de identidad depende de varias características que a menudo entran en conflicto, como las garantías de privacidad y seguridad, la interoperabilidad o la facilidad de uso. En particular, un gran obstáculo en estos sistemas es el propio proveedor de identidad, ya que se convierte en un punto crítico de fallo tanto para la seguridad como para la privacidad. Por ejemplo, un proveedor malintencionado o comprometido puede usurpar o falsificar identidades. Además, puede menoscabar la privacidad de los usuarios mediante el rastreo de su actividad online. El campo de las p-ABCs aplicadas a la gestión de identidad también plantea varios desafíos. Uno de ellos es conseguir la eficiencia necesaria para justificar su uso frente a otras técnicas más sencillas. Del mismo modo, las p-ABC han sufrido una importante falta de estandarización que dificulta su integración en los sistemas del mundo real. Además, intentar conseguir sistemas con funcionalidades avanzadas como pseudónimos o predicados sobre los valores de los atributos, que son uno de los principales atractivos de las p-ABC, agrava aún más estos retos.

Por otro lado, los retos de la gestión de identidad no se limitan a las autenticaciones online tradicionales. A saber, el control de acceso físico requiere demostrar información sobre una persona concreta al tiempo que se protege su derecho a la intimidad. Un ejemplo reciente con repercusión considerable ha sido el “Green Pass” durante la pandemia COVID-19. Los objetivos de privacidad (como ocultar la identidad de la persona concreta, o si estaba vacunada o se había recuperado recientemente) contrastan con la necesidad de seguridad y escalabilidad del proceso de verificación, convirtiendo el problema en un claro desafío. La gestión de identidad en escenarios IoT presenta una serie de retos específicos. Los dispositivos en entornos IoT presentan diversas características, como su elevado número, heterogeneidad o restricciones en capacidad de computación. Los sistemas de identidad deben ser lo suficientemente eficientes y flexibles para adaptarse a ellos. Además, los dispositivos pasan por diferentes fases en su ciclo de vida que deben considerarse más allá de la aplicación de herramientas técnicas para lograr soluciones integrales. Así pues, es necesario abordar de forma coherente diversos desafíos, como la provisión de una raíz de confianza para la identidad de los dispositivos, o de medios para autenticación con garantías de privacidad dentro de un dominio de seguridad.

## 2 Objetivos y metodología

Los diversos retos expuestos dan lugar al objetivo principal que enmarca el desarrollo de esta tesis doctoral: *El diseño de soluciones de manejo de identidad digital seguras, utilizables y que preserven la privacidad tanto de individuos como dispositivos, aplicándolas para abordar retos del mundo real.* Para alcanzar esta meta, se han definido varios objetivos específicos:

**Objetivo 1:** Analizar soluciones de manejo de identidad del estado del arte y sus principales limitaciones.

- Objetivo 2:** Diseñar, desarrollar y evaluar soluciones criptográficas dedicadas a la preservación de la privacidad, especialmente en relación con las credenciales basadas en atributos.
- Objetivo 3:** Especificar métodos y *frameworks* que permitan una gestión de identidad eficaz, utilizable y comprensiva a través de las credenciales basadas en atributos.
- Objetivo 4:** Integrar las soluciones desarrolladas con especificaciones y tecnologías emergentes como los registros distribuidos (DLT) o las credenciales verificables (Verifiable Credentials), evaluando el impacto bidireccional.
- Objetivo 5:** Estudiar, diseñar e implementar la extensión de las nociones de gestión de identidad con preservación de la privacidad a entornos IoT que abarquen el ciclo de vida completo de los dispositivos en función de retos específicos como la heterogeneidad o las limitaciones de recursos.
- Objetivo 6:** Validar las soluciones desarrolladas en diferentes casos de uso y escenarios reales, demostrando su viabilidad.

Estos objetivos han guiado la metodología del trabajo, de carácter iterativo e incremental. En cada ciclo se ha realizado un análisis del estado del arte y el diseño, implementación y evaluación de soluciones técnicas. Además, los resultados han sido continuamente validados en casos de uso de relevancia práctica, particularmente en el contexto de tres proyectos europeos de Horizonte 2020: OLYMPUS [12], CyberSec4Europe [13] y ERATOSTHENES [14]. En el desarrollo de la tesis, se han estudiado y aplicado múltiples tecnologías. Entre ellas destacan las p-ABC, en particular el esquema basado en multi-firmas Pointcheval-Sanders introducido por Camenisch *et al.* [15]. Otra tecnología relevante por sus características de descentralización ha sido los registros distribuidos, como soporte del marco de confianza de ecosistemas de identidad. Del mismo modo, las especificaciones vinculadas a la identidad autosoberana, como las credenciales verificables del W3C, han sido parte del foco por su impacto en las tendencias actuales en materia de identidad. Por último, otras tecnologías como los documentos MUD (*Manufacturer Usage Description*)[16] han desempeñado un papel complementario a las técnicas de identidad investigadas.

### 3 Resultados

La consecución de los objetivos marcados ha resultado en diversas publicaciones en revistas indexadas y conferencias, recogidas en el apartado *Publications* tras la bibliografía. El Cuadro 1 resume los principales resultados obtenidos, así como su relación con los objetivos planteados y dichas publicaciones. Las cuatro publicaciones que conforman el compendio de esta tesis recogen los detalles de estos resultados, que se resumen a continuación.

Resultado	Objetivos	Publicaciones
R0. Análisis de los retos, características y deficiencias de los sistemas de gestión de identidad actuales, particularmente en el ámbito de la privacidad.	1, 5	[17] [18] [19] [20] [21]
R1. Implementación y validación de un esquema de gestión de identidad distribuido con altas garantías de privacidad a través de p-ABCs con divulgación selectiva y no vinculables.	2, 3	[18]
R2. Extensión de la solución de manejo de identidad mediante técnicas de registro distribuido (DLT) y control de acceso basado en confianza-cero para una gestión integral de la confianza y la autorización.	3, 4	[17] [22] [19]
R3. Diseño, implementación y validación de un esquema p-ABC mejorado que cubre características avanzadas de privacidad y seguridad de forma modular y extensible.	2	[21]
R4. Formalización, diseño e instanciación práctica de p-ABCs intransferibles mediante biometría aplicado al control de acceso físico.	2	[23]
R5. Desarrollo de implementaciones de soluciones de mejora de la privacidad adaptadas a dispositivos IoT, especialmente en términos de eficiencia y flexibilidad.	2, 5	[19] [23] [21]
R6. Integración de las soluciones desarrolladas con técnicas estándar de autosoberanía, como la especificación de credenciales verificables, permitiendo aplicaciones interoperables, flexibles y que preservan la privacidad.	4	[17] [20] [19]
R7. Diseño e instanciación de una arquitectura para la gestión de identidad con preservación de la privacidad de dispositivos IoT que cubre los diversos retos específicos identificados.	3, 5	[19] [24]
R8. Validación y evaluación de las soluciones propuestas en diversos escenarios de manejo de identidad para verificar su viabilidad.	6	[17] [25] [26] [27] [18] [19] [24]

Cuadro 1: Resumen de los resultados de la tesis y su relación con los objetivos conseguidos y publicaciones realizadas

### 3.1 Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures

La primera publicación [18] describe la implementación y evaluación del esquema de credenciales basadas en atributos con emisión distribuida (dp-ABC) que actúa como piedra angular del desarrollo de esta tesis. La capacidad de emisión distribuida viene dada por el uso de multi-firmas [15], permitiendo mitigar el problema de punto único de fallo del emisor sin necesidad de depender de complejos procesos de configuración o estrechas relaciones de confianza entre los emisores, lo cual no era posible en la práctica con trabajos existentes en la literatura (R1). En este trabajo, se especifica la aplicación de esta implementación para conseguir un sistema de gestión de identidades distribuido, práctico

y fácil de usar en el marco del proyecto H2020 OLYMPUS (R8). Este sistema, diseñado en búsqueda de usabilidad similar a la de SSO, define un “proveedor de identidad virtual” transparente compuesto por proveedores parciales. Junto con las características de divulgación selectiva y no vinculabilidad aportadas por el esquema dp-ABC, esto permite alcanzar privacidad desde el diseño en el sistema de identidad federada.

El desempeño de la solución fue exhaustivamente evaluado mediante experimentos con la librería desarrollada. El artículo detalla la eficiencia de los procesos del esquema dp-ABC, así como su comportamiento y escalabilidad respecto a los parámetros relevantes (atributos, número de emisores. . .). Además, los resultados se compararon con Idemix [28–30], el sistema de identidad basado en p-ABC que ha recibido un uso más extendido. Nuestra implementación es entre dos y cuatro veces más rápida en la ejecución de todos los procesos de emisión y presentación de credenciales. De hecho, estos resultados se alcanzan con un nivel de seguridad significativamente mayor, 134 bits frente a los 112 bits del método de Idemix. En definitiva, los resultados validan la utilidad práctica de la implementación desarrollada (R1).

### 3.2 Beyond Selective Disclosure: Extending Distributed p-ABC Implementations by Commit-and-Prove Techniques

La segunda publicación [21] se centra en extender la funcionalidad de la implementación de dp-ABC con capacidades de privacidad y seguridad avanzadas. Para ello, modificamos la fase de presentación de la credencial de forma que la prueba de conocimiento cero queda dividida en varias sub-pruebas. Una de las sub-pruebas se dedica a la demostración de posesión del credencial, pudiendo divulgar selectivamente sus atributos o, de forma esencial, ligarlos a un compromiso (*commitment*) de Pedersen [31]. El resto se dedican a probar predicados específicos sobre los atributos, mediante pruebas *commit-and-prove* (comprometer-y-probar). De esta forma, conseguimos la extensibilidad del sistema de forma modular.

Esta mejora se usa para desarrollar un sistema que ofrece pruebas de rango sobre atributos numéricos, mejorando la granularidad en el revelado mínimo de información. También añade soporte para pseudónimos, que se utilizan para habilitar vinculación selectiva controlada por el usuario, y las técnicas de seguridad de inspección y revocación. Para cada una de estas aplicaciones, se midió la sobrecarga de cómputo y uso de memoria. El tiempo de ejecución añadido no sobrepasa los 30 milisegundos por operación, excepto en el caso de las costosas pruebas de rango. Incluso estas últimas cubren la mayoría de casos de uso en un tiempo total inferior a un segundo. Estos resultados validan el interés práctico de la solución. Aún más, el artículo ahonda en el estudio de la viabilidad de la solución con una comparativa con Idemix (en la que nuestra implementación se muestra claramente superior) y resultados favorecedores en entornos móviles y de capacidad reducida (R5).

En definitiva, obtenemos una implementación extensible, usable y práctica de dp-ABCs con funcionalidad avanzada (R3). Estas características son muy deseables en el panorama emergente de casos de uso de identidad basados en SSI, como queda plasmado en la publicación (R0). Así, esta solución puede convertirse en una base para introducir propiedades

avanzadas de privacidad en los escenarios heterogéneos actuales, con una forma clara y sencilla de adaptarla y mejorarla siempre que sea necesario.

### 3.3 To Pass or Not to Pass: Privacy-Preserving Physical Access Control

La tercera publicación [23] presenta la formalización rigurosa e instanciación práctica de las credenciales basadas en atributos ligadas a patrones biométricos (*Biometric-Bound Attribute-Based Credentials*, bb-ABC) (R4). Este marco formal aborda el problema de la transferabilidad de credenciales garantizando que únicamente la persona establecida puede usarlas mediante pruebas biométricas. En particular, la credencial se asocia a un patrón biométrico y el usuario debe demostrar que su característica biométrica medida en el momento de la autenticación coincide (en el sentido de cotejo biométrico) con este patrón. Para ello, se introduce como un nuevo actor en el proceso el lector biométrico como un dispositivo semi-confiable desplegado en las instalaciones del verificador. Para justificar la confianza necesaria en él, su funcionalidad es reducida y reutilizable para cualquier verificador, facilitando procesos de certificación y auditoría en escenarios prácticos. Además, el dispositivo solo requerirá comunicación limitada y material criptográfico efímero, disminuyendo la superficie de ataque.

El artículo detalla la primera formalización de las características de seguridad de completitud, solvencia y no vinculabilidad en este escenario. Además, establece dos construcciones genéricas que se demuestran seguras en este marco formal. La primera, fácilmente instanciable de forma eficiente, delega en el dispositivo lector el proceso de cotejo biométrico entre el nuevo patrón y el ligado a la credencial. La segunda requiere que el usuario pruebe en conocimiento cero que ambos patrones biométricos coinciden. Así, el comportamiento del dispositivo y el verificador puede ser auditado por los usuarios, y demostrado fehacientemente a otras partes. A cambio, la eficiencia de este método está ligada al proceso de cotejo biométrico.

Aún así, el artículo describe una instanciación basada en el esquema dp-ABC, compromisos de Pedersen, AES-GCM-256 y el método de comparación biométrica basada en imágenes faciales de Oumane *et al.* [32]. La seguridad de este método se deriva de las primitivas utilizadas. De hecho, aunque la solvencia del método no se mantendría, la privacidad de las pruebas generadas seguiría manteniéndose incluso en escenarios postcuánticos. La aplicabilidad práctica de la solución se demuestra mediante *micro-benchmarks* de los sobrecostes criptográficos introducidos, utilizando dispositivos relevantes: un ordenador de uso general para el verificador, un móvil para el usuario y una Raspberry 3 como dispositivo lector (R5). A pesar de que la implementación deja lugar a diversas optimizaciones avanzadas, el sobrecoste total está en unos 3 segundos, que se reduce a 2,1 segundos si aplicamos las pre-computaciones básicas. Estos resultados son adecuados para procesos de control de acceso físico seguro, escalable y con privacidad, como se requiere, por ejemplo, en el caso del "Green Pass.<sup>en</sup> pandemias. En términos de comunicación, los resultados tampoco son prohibitivos, requiriendo el intercambio de tres mensajes de tamaño inferior a 135kBs.

### 3.4 A privacy-preserving attribute-based framework for IoT identity lifecycle management

La cuarta publicación [19] identifica los desafíos relacionados con la identidad a lo largo del ciclo de vida de los dispositivos IoT (R0). Desde este punto de partida, definimos una arquitectura y flujos que cubre de forma comprehensiva y coherente cada uno de estos desafíos, desde el provisionamiento de una raíz de confianza en la identidad del dispositivo, hasta la autenticación y autorización de dispositivos con garantías de privacidad siguiendo un modelo de confianza-cero (R2) durante la fase operacional. Esta arquitectura se instanció en el marco del proyecto H2020 ERATOSTHENES (R8). Los componentes utilizados, como documentos MUD o registros distribuidos, habilitan la seguridad y utilidad práctica de la instanciación. Además, la flexibilidad es uno de los principales objetivos, y se prevén variaciones sobre la instanciación básica, que pueden ser acomodadas de forma simple gracias a las pruebas de identidad y el enfoque de confianza-cero (por ejemplo, cambiando el nivel de confianza en dispositivos según sus características de seguridad) (R7).

Entre estos componentes, destaca el uso del esquema dp-ABC durante los procesos de autenticación y autorización. En contraste con los escasos trabajos similares en la literatura, el artículo demuestra la viabilidad práctica de esta solución mediante resultados de rendimiento de la implementación propuesta. Sin embargo, su aplicabilidad en entornos IoT es especialmente propiciada por el enfoque flexible de la solución, que permite el uso de diversos esquemas con características de seguridad, privacidad y eficiencia variables de forma transparente a los flujos de identidad (R5). Este aspecto está alineado con las tendencias de uso de perfiles y objetivos de seguridad [33] en ecosistemas IoT.

La falta de una solución interoperable de p-ABCs como la descrita ha sido uno de los grandes problemas de esta tecnología en el salto a su adopción en la práctica. En gran medida, esto se ha debido a los escasos esfuerzos de estandarización en este entorno. La especificación de credenciales verificables (*Verifiable Credentials*, VC) de W3C detalla un modelo para representar y manejar credenciales digitales, y es una de las tecnologías más representativas de las propuestas de identidad auto-soberana. Sin embargo, las instanciaciones actuales del esquema presentan problemas de vinculabilidad y dificultad para alcanzar divulgación mínima. El artículo define dos *suites* de firma que aglutinan los procesos del esquema dp-ABC, alcanzando una integración transparente con la especificación que ofrece soporte para las funcionalidades avanzadas del esquema. Así, allanamos el camino a alternativas con mejores características de privacidad que las soluciones existentes basadas en VC, y conseguimos la flexibilidad necesaria para aplicar el esquema a entornos heterogéneos (R6).

## 4 Conclusiones y trabajo futuro

El auge generalizado de la digitalización de servicios y la conectividad de usuarios y dispositivos IoT plantea grandes retos en el manejo de la identidad digital. Las soluciones existentes no han conseguido abordar satisfactoriamente lagunas en materia de seguridad y privacidad como que el proveedor de identidad se convierta en un único punto de fallo,

y presentan problemas que dificultan su adopción en escenarios prácticos, como escasa eficiencia, extensibilidad o interoperabilidad. Esta tesis doctoral se ha realizado con el objetivo de desarrollar soluciones de gestión de identidad que sean seguras, provean garantías de privacidad y sean prácticas, llenando los vacíos encontrados en la literatura. En su consecución, las principales conclusiones obtenidas se pueden resumir en:

- Las credenciales basadas en atributos con emisión distribuida (dp-ABC) pueden aplicarse para abordar la gestión de identidad de forma eficiente y utilizable, mitigando al mismo tiempo el problema del punto único de fallo del proveedor de identidad.
- El uso de dp-ABC favorece alcanzar nociones de privacidad, seguridad y funcionalidad inherentes a los ecosistemas basados en la autorización de confianza-cero.
- La modificación del proceso de presentación de las p-ABC, basada en ligar atributos a *compromisos* y el uso de técnicas del tipo "comprometer-y-probar", permite la extensibilidad modular del esquema conservando las garantías formales de seguridad y privacidad.
- En particular, este resultado se utilizó para alcanzar de forma eficiente características avanzadas de p-ABCs relevantes en casos de uso contemporáneos, a saber, inspección, pseudónimos, revocación y pruebas de rango numérico.
- El concepto de bb-ABC, que permite la intransferibilidad de credenciales mediante biometría sin dispositivos dedicados por usuario, se formaliza mediante la definición de variantes de las propiedades de seguridad tradicionales de las p-ABC: corrección, solvencia e inconnexibilidad.
- La relevancia práctica de las bb-ABCs se demostró a través de la definición de dos construcciones complementarias y su instanciación con primitivas concretas, mostrando su eficiencia a través de *micro-benchmarks*.
- La aplicabilidad de las dp-ABC en casos de uso relevantes, y en particular en entornos IoT, mejoró con su integración efectiva en la especificación de *Verifiable Credentials* de W3C.
- Es fundamental contar con un *framework* para la gestión de identidad de los dispositivos IoT a lo largo de su ciclo de vida que sea flexible, comprensivo y plantee garantías de privacidad. Las dp-ABC pueden ser una piedra angular para su consecución.
- Las soluciones desarrolladas se aplicaron con éxito para resolver retos en múltiples casos de uso relevantes, demostrando su relevancia en el panorama actual de identidad digital.

Los resultados de esta tesis doctoral abren el camino a diversas vías futuras de trabajo. En primer lugar, la extensibilidad del esquema dp-ABC desarrollado propicia el estudio de mejoras innovadoras. En esta dirección, una meta especialmente trascendente es ligar

los resultados con la iniciativa de Identidad Digital Europea basada en el uso de carteras digitales. Nuestros resultados pueden servir de base para extender las capacidades ofrecidas por los despliegues iniciales, potenciando el control de los ciudadanos sobre su derecho a la privacidad.

Por otro lado, el *framework* genérico de gestión de identidad IoT desarrollado puede ser instanciado a través de otras herramientas de última generación y adaptarse a escenarios específicos. En este sentido, se puede ahondar en esta temática para cubrir las necesidades de dispositivos de gama ultrabaja, por ejemplo, mediante primitivas de criptografía simétrica. En esta dirección, el avance en sistemas de perfiles y objetivos de seguridad se presenta como un paso clave para alcanzar enfoques flexibles de gestión de identidad y ciberseguridad en general.

Por último, siguiendo la tendencia general en el mundo de la criptografía, es posible explorar la consecución de nociones de seguridad y privacidad equivalentes en el mundo post-cuántico. En particular, siguiendo la tendencia de flexibilidad y configurabilidad de las soluciones desarrolladas en la tesis, se distingue una línea de investigación e innovación para allanar la transición entre la criptografía tradicional y las soluciones post-cuánticas mediante su aplicación complementaria.



# Abstract

## 1 Motivation

Digital services have become a basic element of our society. In this environment, it is necessary to authenticate the various digital entities involved, and in particular the users. This requirement is not only related to identification, but also to the authentication of identity attributes such as e-mail or a person's age. Identity management solutions must therefore be implemented to achieve a secure ecosystem. The ubiquitous nature of the need for identity management makes the security and privacy properties achieved by these solutions more relevant than ever.

Traditional identity systems, which see widespread adoption due to their ease of use, produce serious security and privacy problems. For example, breaches of service provider databases have led to the leakage of billions of records containing user data such as emails, passwords and personal information [1–3]. Single sign-on (SSO) systems mitigate the problem by reusing authentication information through third party identity providers (IdPs) to access multiple services. However, this leads to single point of failure concerns regarding security and privacy, culminating in the exhaustive tracking of all users' online activity. In fact, these mechanisms are generally offered by large technology companies such as Google or Facebook, which have already been under fire for their misuse of users' data [4].

These security and privacy issues are not limited to traditional user interactions with digital services. As the European Economic and Social Committee (EESC) stresses, the evolution of digital technologies and their application to the digitisation of society offers enormous opportunities for economic growth, justice and inclusion, but it is not without threats. Risks such as breaches of privacy or the introduction of prejudice and discrimination increase considerably. This is why European policies and legislation are now moving towards digital sovereignty.<sup>2</sup>

The rise of the Internet of Things (IoT), where machines or devices are the main actors in the exchange of information, is one of the main exponents of this trend towards digitisation. In 2023, the number of connected IoT devices reached 15 billion, and estimates for the coming years reflect significant growth, reaching 30 billion devices by 2030 [6]. These devices have highly variable characteristics in terms of computing capacity, memory or power consumption limitations. Their ubiquity and the sensitive data shared in IoT envi-

---

<sup>2</sup><https://www.eesc.europa.eu/en/our-work/opinions-information-reports/opinions/digital-identity-data-sovereignty-and-path-towards-just-digital-transition-citizens-living-information-society>

ronments emphasise the need for secure solutions with privacy mechanisms, particularly in the area of identity management.

In recent years, the trend has shifted towards more privacy-friendly solutions that place users (and their devices) at the centre of identity management systems. In this sense, the emerging paradigm of self-sovereignty in identity management (SSI) [7] establishes precepts such as data minimisation, non-traceability, or user consent. The importance of these concepts is highlighted in regulations such as the GDPR. Particularly, the European Commission aims to make digital identity available to all EU citizens through the eIDAS2 regulation [8]. This proposal establishes the use of identity wallets for authentication in digital services across the EU following the principles of self-sovereignty. On the other hand, the SSI approach is not only suitable for individuals, but is also applicable to the IoT world and, in general, to any digital service or object [9].

In the pursuit of these ambitious privacy goals, the need to apply specifically designed cryptographic tools, known as privacy enhancing technologies (PETs), has become apparent. PETs make it possible to collect, process, analyse or share information while protecting the confidentiality of personal data. They aim for a variety of purposes such as obfuscating data (through manipulation or hiding), enabling the processing of encrypted data, or performing distributed data analysis. These techniques cannot be considered a "silver-bullet" that solves all privacy and data protection problems and their implementation still presents some barriers, such as their reduced accessibility to policy-makers due to their innovative nature. Nevertheless, they are an important tool to achieve the paradigm of privacy by design. In addition, they enable new applications and use cases, as they allow the guarantee of privacy rights while maintaining data usability. Among existing PETs, privacy-preserving Attribute-Based Credentials (p-ABC) [10, 11] stand out for their suitability for general-purpose identity management, empowering users' control over their online identity. In fact, they are considered a key element for future privacy-enhancing developments in the European digital identity solution. Thus, in this thesis we focus on p-ABCs and their application to solve identity management challenges.

P-ABCs allow the issuance of digital credentials that certify the user's attributes. Once in the user's possession, these credentials can be used without third party intervention to derive unlinkable *tokens* through zero-knowledge proofs, selectively disclosing some of the information contained therein. In this process, the verifier will still have formal assurances that the disclosed data was actually certified by the issuer. Thus, p-ABCs can be applied to perform authentication of identity attributes related to fine-grained authorisation processes while maintaining high privacy guarantees against the service provider. In addition, the problem of user tracking by central entities such as the identity provider is mitigated.

Despite previous efforts, the field of identity management still poses challenges in terms of security and privacy preservation. First of all, the applicability of identity management systems depends on several often conflicting characteristics, such as privacy and security guarantees, interoperability or ease of use. In particular, a major obstacle in these systems is the identity provider itself, as it becomes a critical point of failure for both security and privacy. For example, a malicious or compromised provider can usurp or forge identities. Also, it can undermine users' privacy by tracking their online activity. The field of p-

ABCs applied to identity management also poses several challenges. One of them is to achieve the necessary efficiency to justify their use compared to other, simpler techniques. Similarly, p-ABCs have suffered from a significant lack of standardisation, which makes it difficult to integrate them into real-world systems. In addition, trying to achieve systems with advanced functionalities such as pseudonyms or predicates on attribute values, which are one of the main attractions of p-ABCs, further exacerbates these challenges.

The challenges of identity management are not limited to traditional online authentication. For instance, physical access control requires proving information about a specific individual while protecting her right to privacy. A recent example with considerable impact is the ‘Green Pass’ during the COVID-19 pandemic. The privacy objectives (e.g., hiding the identity of the specific person, or whether they were vaccinated or recently recovered) contrast with the need for security and scalability of the verification process, posing a challenging scenario. On another note, identity management in IoT ecosystems presents several specific challenges, derived from devices’ characteristics, such as their high numbers, heterogeneity or computational capacity constraints. Identity systems must be efficient and flexible enough to adapt to them. In addition, devices go through different phases in their lifecycle that need to be considered beyond the application of technical tools to achieve holistic solutions. Thus, various challenges, such as the provision of a root of trust for device identity or means for privacy-preserving authentication within a security domain, need to be addressed in a coherent manner.

## 2 Objectives and Methodology

From the various challenges outlined above, the main objective that frames the development of this doctoral thesis is raised: *The design of secure, usable and privacy-preserving digital identity management solutions for both individuals and devices, applying them to address real-world challenges.* To achieve this goal, it has been divided into six specific objectives:

- Objective 1:** Analyse state of art identity management solutions and their main limitations.
- Objective 2:** Design, develop and evaluate cryptographic solutions devoted to privacy preservation, particularly related to attribute-based credentials.
- Objective 3:** Specify novel methods and frameworks for enabling efficient, usable and comprehensive identity management based on attribute-based credentials.
- Objective 4:** Integrate developed solutions with emerging specifications and technologies such as distributed ledger technologies or verifiable credentials, evaluating the bidirectional impact.
- Objective 5:** Study, design and implement the extension of privacy-preserving identity management notions to IoT environments covering devices’ complete life-

cycle according to specific challenges such as heterogeneity or resource constraints.

**Objective 6:** Validate the developed solutions over different use cases and real scenarios, demonstrating their practicality.

These objectives have guided the iterative and incremental methodology of the work. In each cycle, we carried out an analysis of the state of the art and the design, implementation and evaluation of technical solutions. In addition, the results have been continuously validated in use cases of practical relevance, particularly in the context of three European Horizon 2020 projects: OLYMPUS [12], CyberSec4Europe [13] y ERATOSTHENES [14]. In the development of the thesis, multiple technologies have been studied and applied, as is detailed in the compendium's publications. Among them, p-ABCs stand out, especially the Pointcheval-Sanders multi-signature scheme introduced in [15]. Another relevant technology because of its decentralisation features has been Distributed Ledger Technologies (DLT), as a supporting tool for the identity ecosystem trust framework. Similarly, specifications linked to self-sovereign identity, such as W3C's Verifiable Credentials, have been part of the focus due to their impact on current digital identity trends. Finally, other technologies have played a complementary role to the identity techniques investigated, as is the case of Manufacturer Usage Description (MUD) files [16] and their strong irruption in zero-touch IoT scenarios.

## 3 Results

The achievement of the objectives set has resulted in various publications in indexed journals and conferences, listed in *Publications*. Table 2 summarises the main results obtained, as well as their relationship with the the set of objectives and the aforementioned publications. The four publications that make up the compendium of this thesis contain the details of these results, which are summarised below.

### 3.1 Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures

The first publication [18] describes the implementation and evaluation of the distributed privacy-preserving attribute-based credential scheme (dp-ABC) that is a cornerstone of the development of this thesis. The distributed issuance capability is provided by the use of multi-signatures, as detailed by the authors of the theoretical definition of the scheme [15]. This enables mitigating the single point of failure problem of the issuer without having to rely on complex configuration processes or trust relationships between issuers, which was not possible in practice with existing works in the literature (R1). In this paper, we apply this implementation to achieve a practical and user-friendly distributed identity management system in the context of the H2020 OLYMPUS project (R8). This system, designed in pursuit of SSO-like usability, defines a transparent "virtual identity

<b>Result</b>	<b>Objectives</b>	<b>Publications</b>
R0. Analysis of the properties, challenges and deficiencies of current identity management systems, particularly focusing on privacy.	1, 5	[17] [18] [19] [20] [21]
R1. Implementation and validation of a distributed privacy-preserving identity management scheme with high privacy guarantees through p-ABCs with selective disclosure and un-linkability.	2, 3	[18]
R2. Enhancement of the identity solution with Distributed Ledger Technologies and zero-trust based access control for comprehensive trust and authorisation management.	3, 4	[17] [22] [19]
R3. Design, implementation and validation of an extended p-ABC scheme covering advanced privacy and security features in a modular and extensible way.	2	[21]
R4. Formalisation, design and instantiation of practical non-transferable p-ABCs through biometrics for physical access control.	2	[23]
R5. Development of implementations of privacy-enhancing solutions suitable for IoT devices, particularly in terms of efficiency and flexibility.	2, 5	[19] [23] [21]
R6. Integration of developed solutions with standard self-sovereign practices such as the Verifiable Credentials specification enabling interoperable, flexible and privacy-preserving applications.	4	[17] [20] [19]
R7. Design and instantiation of a framework for privacy-preserving identity management of IoT devices covering identified challenges.	3, 5	[19] [24]
R8. Validation and evaluation of the proposed solutions in different identity management scenarios in order to verify their feasibility.	6	[17] [25] [26] [27] [18] [19] [24]

Table 2: Summary of results and their link to objectives and publications achieved

provider" composed of partial providers. Together with the selective disclosure and un-linkability features provided by the dp-ABC scheme, this allows for achieving privacy by design in a federated identity system.

The performance of the solution was thoroughly evaluated through experiments with the developed library. The paper details the efficiency of the processes of the dp-ABC scheme, as well as their behaviour and scalability with respect to the relevant parameters (attributes, number of issuers...). Furthermore, the results were compared with Idemix [28–30], the p-ABC-based identity system that has received the most widespread use. Our implementation is two to four times faster in the execution of all credential issuance and presentation processes. Moreover, these results are achieved with a significantly higher level of security, 134 bits versus the 112 bits for the Idemix configuration used. In short, the results validate the practical utility of the developed implementation

(R1).

### 3.2 Beyond Selective Disclosure: Extending Distributed p-ABC Implementations by Commit-and-Prove Techniques

The second publication [21] focuses on extending the functionality of the dp-ABC implementation with advanced privacy and security capabilities. To this end, we modify the credential presentation phase in such a way that the zero-knowledge test is split into several sub-tests. One of the sub-tests is dedicated to the proof of possession of the credential, being able to selectively disclose its attributes or, crucially, to bind them to a Pedersen commitment [31]. The rest are devoted to proving specific predicates on the attributes through commit-and-prove proofs. In this way, we achieve extensibility of the system in a modular way.

This enhancement was used to develop a system that supports range proofs on numeric attributes, improving granularity in the minimal disclosure of information. We also added support for pseudonyms, which are used to enable user-controlled selective linkability, and the security techniques of inspection and revocation. The computational and memory overhead with respect to the basic scheme was measured for each of these applications. The added execution time does not exceed 30 milliseconds per operation, except in the case of range proofs. Even the latter, which are generally expensive, can cover most use cases in a total time of less than one second. These results validate the practical relevance of the solution. Furthermore, the paper delves deeper into the feasibility study of the solution with a comparison with Idemix (where our implementation shows superior performance) and favourable results in mobile and low-capacity environments (R5). Ultimately, we obtain an extensible, usable and practical implementation of dp-ABCs with advanced functionality (R3). These features are highly desirable in the emerging landscape of SSI-based identity use cases, as detailed in the publication (R0). Thus, the solution can become a basis for introducing advanced privacy properties in today's heterogeneous scenarios, with a clear and simple way to adapt and enhance it whenever necessary.

### 3.3 To Pass or Not to Pass: Privacy-Preserving Physical Access Control

The third publication [23] presents the rigorous formalisation and practical instantiation of biometric-bound attribute-based credentials (bb-ABCs) (R4). This framework addresses the problem of credential transferability by ensuring that only the established person can use the credential through biometric matching. In particular, the credential is associated to a biometric template and the user must prove that her biometric characteristic measured at the time of authentication matches this template. For this purpose, the biometric reader is introduced in the process, acting as a semi-trusted device deployed at the verifier's premises. To justify the necessary reliance on the reader, its functionality is reduced and reusable for any verifier, facilitating certification and auditing processes in

practical scenarios. Furthermore, the device will only require limited communication and ephemeral cryptographic material, decreasing the attack surface.

The paper details the first formalisation of the security properties of correctness, soundness and unlinkability in this scenario. Additionally, it establishes two generic constructs that are proven secure within this framework. The first, easily instantiated in an efficient way, delegates to the reader device the biometric matching process between the new template and the one bound to the credential. The second requires the user to prove in zero knowledge that both biometric templates match. In this way, the behaviour of the device and the verifier can be audited by users, and reliably demonstrated to other parties. In return, the efficiency of this method is tied to the biometric matching process.

Even so, the article describes an instantiation based on the dp-ABC scheme, Pedersen commitments, AES-GCM-256 and Oumane's facial image-based biometric matching method *et al.* [32]. The security of this method is assured by the primitives used. In fact, although soundness would not hold, the privacy of the generated proofs is still maintained even in post-quantum scenarios.

The practical applicability of the solution is demonstrated through micro-benchmarks of the cryptographic overheads introduced, using relevant devices: a general-purpose laptop for the verifier, a mobile phone for the user and a Raspberry 3 as a reader device (R5). Although the implementation leaves room for several advanced optimisations, the total overhead is around 3 seconds, which is reduced to 2.1 seconds if we apply basic pre-computations. These results are suitable for secure, scalable and privacy-preserving physical access control processes, as required, for example, in the case of "Green Pass" in pandemics. In terms of communication, the results are also not prohibitive, requiring the exchange of three messages smaller than 135kBs.

### 3.4 A privacy-preserving attribute-based framework for IoT identity lifecycle management

The fourth publication [19] identifies identity-related challenges throughout an IoT device lifecycle (R0). From this starting point, we defined an architecture and flows that comprehensively and coherently cover each of these challenges, from the provisioning of a root of trust for the device identity, to the privacy-preserving authentication and authorisation of devices following a zero-trust approach during the operational phase (R2). This framework was instantiated in the context of the H2020 project ERATOSTHENES (R8). The components used, such as MUD files or DLTs, enable the security and practical usability of the instantiation. In addition, flexibility is one of the main objectives, and variations on the basic instantiation are foreseen. This can be accommodated in a simple way thanks to the identity proofing and trust-zero approach (e.g., changing the trust level on devices according to their security characteristics) (R7).

Among these components, the use of the dp-ABC scheme during the authentication and authorisation processes stands out. In contrast to the few similar works in the literature, the paper demonstrates the practical feasibility of this solution through performance results of the proposed implementation. Nevertheless, its applicability in IoT environments is especially favoured by the flexible approach of the solution, which allows the

use of several primitives with varying security, privacy and efficiency characteristics in a transparent way to the identity flows (R5). This aspect is aligned with trends in the use of security profiles and targets in IoT ecosystems [33].

The lack of an interoperable p-ABC solution such as described above has been one of the major issues of this technology in the leap to practical adoption. To a large extent, this has been due to the limited standardisation efforts in this scope. The W3C's Verifiable Credentials (VC) specification details a model for representing and managing digital credentials, and is one of the most representative technologies for SSI approaches. However, current instantiations of the specification present problems of linkability and difficulty in achieving minimum disclosure. The paper defines two signature suites that agglutinate the processes of the dp-ABC scheme, achieving a seamless integration with the specification that provides support for the advanced functionalities of the scheme. Thus, we pave the way for alternatives with better privacy characteristics than existing VC-based solutions, and we achieve the necessary flexibility to apply the scheme to heterogeneous environments (R6).

## 4 Conclusions and future work

The widespread rise of digitisation of services and connectivity of users and IoT devices poses major challenges for digital identity management. Existing solutions have not been able to satisfactorily address security and privacy gaps such as the identity provider becoming a single point of failure, and present problems that hinder their adoption in practical scenarios, such as low efficiency, extensibility or interoperability. This thesis has been carried out with the aim of developing identity management solutions that are secure, practical and provide privacy guarantees, filling the gaps found in the literature. In its achievement, the main conclusions obtained can be summarised as follows:

- Distributed privacy-preserving Attribute-Based Credentials (dp-ABC) have been applied to address identity management in an efficient and usable way while mitigating the single point of failure issue of the issuer.
- The use of dp-ABC enables privacy, security and functionality notions inherent to ecosystems based on zero-trust authorisation.
- The leveraged presentation process of p-ABCs based on linking commitments and commit-and-prove techniques enables modular extensibility retaining formal guarantees of security and privacy.
- The p-ABC extension was demonstrated as a mechanism to efficiently achieve key advanced p-ABC features for contemporary identity use cases, namely inspection, pseudonyms, revocation and range proofs.
- The concept of Biometric-Bound Attribute-Based Credentials (bb-ABC), which enables non-transferability through biometrics without dedicated devices per-user, has been formalized through the definition of variants of the traditional p-ABC security properties: correctness, soundness and unlinkability.



- The practical relevance of bb-ABC was demonstrated through two complementary constructions and their instantiation with concrete primitives, whose efficiency was showcased through micro-benchmarks.
- The applicability of dp-ABC in relevant use cases, particularly IoT environments, was improved through their effective integration into the W3C Verifiable Credential specification.
- A framework for flexible, comprehensive and privacy-preserving identity management of IoT devices throughout their lifecycle is crucial, and has been achieved with dp-ABCs as a key enabler.
- The developed solutions were successfully applied to solve challenges in multiple relevant use cases, showcasing their relevance in the current identity landscape.

The outcomes of this thesis open the way to several avenues of future work. Firstly, the extensibility of the developed dp-ABC scheme allows for innovative improvements. In this direction, a particularly important goal is to link the results to the European Digital Identity initiative based on the use of digital wallets. Our results can serve as a basis for extending the capabilities offered by the initial deployments, enhancing citizens' control over their right to privacy and services offered by identity providers.

On the other hand, the generic IoT identity management framework developed can be instantiated through other state-of-the-art tools and adapted to specific scenarios. In this sense, it can be further developed to cover the needs of ultra-low-end devices, for example, by means of symmetric cryptographic primitives. In this direction, the advancement of systems based on security profiles and targets is a key step towards flexible approaches to identity management and cyber-security in general.

Finally, following the general trend in the world of cryptography, one may explore the achievement of equivalent notions of security and privacy in the post-quantum world. In particular, following the trend of flexibility and configurability of the solutions developed in the thesis, a line of research and innovation can be the smoothing of the transition between traditional cryptography and post-quantum solutions through their complementary application.

# Chapter 1

## Introduction

Digital services have become a staple of our society. Online presence of users and digitisation of services are prevalent, and ever on the rise with trending shifts towards “as-a-service” paradigms. In this context, it is necessary to authenticate digital entities, and particularly users. This need does not only encompass digital identification, but also the authentication of identity attributes such as a person’s e-mail or age. For instance, it is typical to offer access to restricted services depending on a condition, such as being older than a threshold age. Thus, identity management solutions must be applied to achieve a safe ecosystem. The ubiquitous nature of this need makes security and privacy properties achieved by these solutions more relevant than ever.

Traditional identity systems, being widespread because of their usability, have led to serious security and privacy issues. For instance, breaches to databases in service providers [1–3] have led to the leakage of billions of records with user information such as e-mails, passwords and personal information. User security and privacy is jeopardised, not only on the affected services but in their general online presence. Single Sign-On (SSO) systems mitigate the issue by reusing authentication information through third party providers to access multiple online services. However, this setting introduces grave single-point of failure issues in terms of security and privacy, with the identity provider holding control over all processes and particularly becoming a huge user tracking entity. Further, these authentication and authorisation mechanisms have often been offered by Big Tech companies like Google or Facebook, which have already been in the spotlight for misusing user data [4].

These security and privacy concerns are not restricted to traditional user interactions with digital services. As highlighted by the European Economic and Social Committee (EESC), the evolution of digital technologies and their application to society’s digitisation brings enormous opportunities for economical growth, justice and inclusion, but is not exempt of key threats. Privacy violations or introduction of biases and discrimination become potential risks, especially with the push to impose products and services by large tech companies. Thus, European policies and legislation are now moving towards digital sovereignty.<sup>1</sup>

---

<sup>1</sup><https://www.eesc.europa.eu/en/our-work/opinions-information-reports/opinions/digital-identity-data-sovereignty-and-path-towards-just-digital-transition-citizens->

The accelerated advent of the Internet of Things (IoT), where machines or devices are the main actors in information sharing [5], is one of the main exponents of this trend towards digitisation. In 2023, the number of connected IoT devices reached 15 billions, and predicted estimates for next years reflect significant growth, reaching 30 billions of devices in 2030 [6].<sup>2</sup> The interconnected devices have highly variable characteristics in terms of computational power, memory, or energy consumption limitations. The pervasiveness of such devices and the sensitive data shared in these environments heightens the significance of pursuing security and privacy in IoT processes. Tailored identity management solutions will be crucial for achieving these goals.

In recent years, there has been a shift towards more privacy-friendly solutions, putting users (and their devices) at the center of identity management systems. This has led to the notion of self-sovereignty in identity management [7], with multiple precepts like data minimisation or user consent. The importance of these concepts has been brought to attention and is supported by regulations like GDPR or the original eIDAS. Particularly, in pursue of digital sovereignty, the European Commission is aiming to make digital identity available to all EU citizens through the eIDAS2 regulation [8]. The proposal intends to establish identity wallets Widely usable for authentication to digital services across the EU following self-sovereign identity (SSI) principles such as giving full control to users to choose which aspects of their identity, data and certificates they share with third parties.<sup>3</sup> The SSI approach is not only apt for individuals, but is also being introduced in the IoT world [9] and in general any digital service and object.

In the pursue of fulfilling these ambitious privacy goals, it has become apparent that it is necessary to apply cryptographic tools with tailored features. The research community has been highly active on developing Privacy-Enhancing Technologies (PET) for many years. PETs enable the collection, processing, analysis or sharing of information while protecting the confidentiality of personal data. They have various purposes, such as obfuscation of data (through manipulation or hiding), enabling processing of encrypted data or distributed analytics. These techniques cannot be considered as a “silver bullet” to solve all privacy and data protection challenges, and their implantation still presents some barriers such as their innovative nature making them poorly accessible to policy makers. Nonetheless, they are a key enabler for achieving the privacy-by-design paradigm. Moreover, they enable new applications and use cases, as privacy rights can be ensured while maintaining *data utility* [34]. Among existing PETs, privacy-preserving Attribute-Based Credentials (p-ABC) [10, 11], also known as anonymous credentials, stand out as fitting for general purpose identity management that gives users control over their online identity. In fact, although not included in initial versions, they are considered as a key enabler for future developments that enhance privacy in the European Digital Identity solution.<sup>4</sup> Thus, in this thesis we focus on p-ABCs and their application to solving key

---

living-information-society

<sup>2</sup>Other estimates vary heavily (reaching numbers close to 100 billion) mainly because of different criteria for defining IoT devices. Nonetheless, all of them agree on high magnitudes and the significant increasing trend.

<sup>3</sup>[https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity\\_en](https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity_en)

<sup>4</sup><https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0281>

identity management challenges.

p-ABCs enable the issuance of digital credentials certifying user attributes. These credentials can then later be used, without input from external parties, to derive unlinkable tokens through zero-knowledge proofs that selectively disclose some of the information contained in the credential. Nonetheless, the verifier will still have formal guarantees that the data revealed was really certified by the issuer. Thus, p-ABCs can be applied to perform the authentication of identity attributes related to fine-grained authorisation process, e.g., when access to a service requires a specific nationality or income level, while keeping high privacy guarantees against the verifier. In fact, even when multiple entities collude, they cannot collect more information than the user intended to reveal during the authentication process. Particularly, this mitigates the issue of user tracking by central entities like the identity provider. Of course, this only holds while no identifying information, such as the unique social security number, is revealed by the user. Within this thesis, we work on creating systems by employing and enhancing p-ABCs with distributed issuance, establishing links to relevant identity and trust frameworks, and ensuring interoperability and efficient application to IoT scenarios.

Despite the efforts and progress, the field of identity management poses several challenges in terms of security and privacy preservation. First, the applicability of identity management systems hinges on them achieving several characteristics that are often conflict with each other, such as privacy and security guarantees, interoperability or user-friendliness. Particularly, a big weakness in these systems is the identity provider itself, as it becomes a key point of failure for both security and privacy. For instance, a malicious or compromised provider may forge identities or incur in identity theft attacks. Also, it may leak user data, or jeopardise user privacy through tracking mechanisms.

The field of privacy-enhancing technologies applied to identity management, and particular p-ABCs, also poses several challenges. One of them is achieving efficient enough solutions to justify their use over other simpler techniques. On a similar note, interoperability fosters adoption, but p-ABCs have suffered a significant lack of standardisation hampering their integration in real-world systems. Additionally, such solutions suffer from usability issues, both from users and developers point of views. Lastly, trying to achieve systems with advanced functionalities such as pseudonyms or predicates over values, which are one of the main appealing properties of p-ABCs in terms of privacy and security, further exacerbates these challenges.

As introduced before, the challenges of identity management do not stop at traditional online authentications. For instance, privacy-preserving physical access control requires proving information about a specific individual while protecting their right to privacy. A recent example that has brought plenty of attention to such an use case has been the “Green Pass” during the COVID-19 pandemic. The privacy goals (e.g., hiding the specific person’s identity, or whether she was vaccinated or recently recovered) in contrast with the security and scalability of the verification process make achieving a satisfactory solution challenging.

On a different note, the identity management in IoT scenarios presents a different set of challenges. Devices in IoT environments present various characteristics such as their high numbers, heterogeneity, or constraints in various aspects like computation or power

consumption. Identity systems must then be efficient and especially flexible enough to accommodate them. Additionally, devices will go through different lifecycle phases in a mostly automated way, which must be covered beyond the application of technical tools for achieving comprehensive solutions. Thus, various challenges such as providing root of trust for device identities, or means for privacy-preserving authentication within a security domain must be addressed in a coherent way.

## 1.1 Objectives

The development of this thesis has been carried out under a main objective, from which several specific objectives were derived. The main objective is detailed below:

Devise, design and implement usable, secure and privacy-preserving solutions for managing individuals and devices digital identities, applying them to address challenges in the real world.

This general objective has been broken down into the following specific objectives:

- Objective 1:** Analyse state of art identity management solutions and their main limitations.
- Objective 2:** Design, develop and evaluate cryptographic solutions devoted to privacy preservation, particularly related to attribute-based credentials.
- Objective 3:** Specify novel methods and frameworks for enabling efficient, usable and comprehensive identity management based on attribute-based credentials.
- Objective 4:** Integrate developed solutions with emerging specifications and technologies such as distributed ledger technologies or verifiable credentials, evaluating the bidirectional impact.
- Objective 5:** Study, design and implement the extension of privacy-preserving identity management notions to IoT environments covering devices' complete lifecycle according to specific challenges such as heterogeneity or resource constraints.
- Objective 6:** Validate the developed solutions over different use cases and real scenarios, demonstrating their practicality.

## 1.2 Methodology

To achieve the proposed goals, the work has been carried out following an iterative and incremental methodology with multiple cycles. Each cycle has consisted on an analysis of the state of art and its gaps in the tackled topic, the design, implementation and

evaluation of solutions addressing those gaps, and the validation of the results in the resolution of practical use cases. Particularly, much of the work has been carried out within the context of three Horizon 2020 European projects, in which the software developed has been integrated, exploited and validated: OLYMPUS [12], CyberSec4Europe [13] y ERATOSTHENES [14].

Throughout the work, multiple technologies have been studied and applied. From those, the main focus of the thesis has been privacy-preserving Attribute-Based Credentials, particularly the recent scheme based on Pointcheval-Sanders multi-signatures introduced in [15]. Another relevant technology has been Distributed Ledgers, particularly smart-contract-enabled Blockchains, as a verifiable data registry to decentralise the trust framework in an identity ecosystem. Similarly, the specifications linked to Self-Sovereign Identity, such as W3C's Verifiable Credentials, have been a focal point because of their impact in current identity trends. Lastly, other technologies have served a complementary role to the researched identity techniques, like Manufacturer Usage Description (MUD)[16] files as a security configuration tool in IoT ecosystems.

The remainder of this document is organised as follows. Chapter 2 details the results derived within this thesis, describing the analysis of the related work, the identified gaps, and how the thesis addresses them through its outcomes, including the compendium of publications. Chapter 3 presents the conclusions derived from the development of the thesis and future avenues of work enabled by the produced results. Chapter 4 includes the summary and complete text of the four publications that are part of the thesis' compendium.

# Chapter 2

## Summary of results

This chapter describes the main results achieved during the development of the thesis. In order to contextualise them, it first gives an overview on the related work found in the literature.

### 2.1 Related work

This section gives an overview on the related work analysed during the development of the thesis. It covers various topics related to the main work in the thesis. Namely, digital identity management systems and their evolution since traditional approaches, privacy-preserving Attribute-Based Credentials as a key tool for enabling privacy-preserving identity management, and specific challenges and solutions for IoT identity management. We summarise existing works and identify gaps in the literature.

#### 2.1.1 Digital identity management

The use of online services inevitably leads to the challenge of identifying, authenticating and in general managing the identity of (digital) entities. This entails the creation of mechanisms so that a verifier (or relying party) can trust the potential user's identity information. A traditional approach to this challenge is the use of digital credentials such as X.509 certificates [35]. These solutions are based on having an issuer (or Certification Authority) sign a certificate formed by a public key associated to the user and other identity attributes. Users can then authenticate by proving possession of the corresponding private key. The trustworthiness of the process is tackled through PKIs [36]. User credentials are issued by certification authorities (CA), and the relying party needs to trust this authority as a valid issuer. To ensure scalable, secure, and efficient discovery of public keys, the PKI establishes hierarchical relationships. Thus, certification authorities are accredited by other CAs, which recursively share the same process until reaching a *root* authority.

This approach, although simple and useful, presents some clear limitations. First, the PKI hierarchy can be difficult to maintain, and is particularly prone to attacks that lead to impersonation or identity forgery. Second, it places a burden on users to manage

certificates and private keys especially across multiple devices, which is a significant risk. Lastly, traditional certificates do not offer features like minimal disclosure or unlinkability, harming user privacy.

Another traditional solution, overwhelmingly widespread in online service access, is the use of authentication based on information known to the user, such as username and password. This approach is generally user-friendly, but also prone to attacks like guessing or phishing, especially when users need to manage accounts for multiple services [1, 3]. There are ways to improve the safety of the solution like using multi-factor authentication. Here, users need to provide further proof of their identity using different kinds of information, e.g. possession of an e-mail or device [37]. Nonetheless, they help with user identification, but do not tackle the issue of attribute verification by themselves.

Federated identity systems establish a central trust point that manages authentication for multiple service providers [38]. This authority is known as the Identity Provider (IdP), and the user only has to authenticate against it, as in SSO. The centralisation avoids the need of introducing the same data multiple times in different registration processes, and is in general user-friendly. Additionally, revocation is trivial because of the central role of the IdP. The authentication against the IdP itself may be carried out in various manners, such as the ones described above. The IdP will be in possession of the means to identify the user, along with her authenticated attributes like name or date of birth. These attributes can be issued by the own IdP, or obtained from other attribute providers during registration.

Federated identity systems have been widely deployed with several implementations. One such implementation is SAML [39], with a business-centric approach based on issuance of tokens by the issuer specified through XML documents. The OAuth [40] standard, with specific implementations like OpenID Connect [41], has been adopted by commonly used federated IdPs like Google, Facebook or Twitter.

Federated systems offer the capability to provide granular authentication. This system enables the IdP to validate only the necessary attributes when accessing a service. For instance, if a service needs to confirm that a user is over 18 years old, the IdP can assert that the user meets the age requirement without revealing the exact date of birth. However, similar to X509, the issue of linkability persists, as user identification information is shared with each service provider, although some solutions mitigate this issue, like OpenID's pairwise identifiers. This linkability problem is exacerbated by traceability, as the IdP learns every user interactions, potentially creating a detailed profile akin to "Big Brother". Additionally, the IdP becomes a single point of failure; if compromised, all the personal information it holds may be at risk of exposure, and identity forgery and theft are direct.

Recently, there has been a notable push for enhancing user control over identity, particularly with the emergence of the SSI model [7]. SSI advocates for principles of privacy (e.g., minimal disclosure), controllability (e.g., through user consent), and portability (e.g., transparency), fostering a paradigm shift. Proposals rooted in SSI offer superior privacy properties compared to traditional identity paradigms. Instead of relying solely on certificates issued by centralised Public-Key Infrastructures, SSI advocates for decentralisation, user control, and fine-grained attribute authentication. The data space



ecosystem, particularly the GAIA-X framework [42], exemplifies this trend. It initially prioritised federated identity but is now shifting towards integrating SSI principles within the technical convergence framework of the European Data Spaces Business Alliance.<sup>1</sup> The European Digital Identity Wallet (EUDIW) initiative under eIDAS2 [8] serves as another notable example of the increasing significance and adoption of such concepts.

While the SSI paradigm is beyond specific technologies, two specifications have become a cornerstone in many solutions based on SSI: W3C’s Verifiable Credentials (VC) [43] and Decentralized identifiers (DID) [44]. The VC specification aims to achieve a machine-verifiable, secure and flexible way of representing digital credentials, particularly those that attributes linked to a subject. Meanwhile, DIDs are identifiers that associate a subject to a *DID Document* that contains information (i.e., public keys) that allows the subject to prove control over it without relying on a central entity. The plain application of these technologies does not equate to achieving SSI principles. In fact, they can lead to privacy issues [20, 45, 46], although solutions like SD-JWT [47], which offers selective disclosure but no unlinkability, are being explored. Nonetheless, their importance is only rising, as exemplified by the extensions to OIDC to make them usable in its widespread flows [48–50] or their tentative use in key initiatives like data spaces [51] or the European Blockchain Services Infrastructure (EBSI).<sup>2</sup>

In this scope of decentralisation, Distributed Ledger Technologies (DLT) are also very relevant because of their good fit as a Verifiable Data Registry (VDR). The primary exponent of DLTs is Blockchain, with solutions like Indy [52] or Ethereum [53] being applied in the identity management realm (see, e.g. [54, 55]). Other solutions have been explored for filling this role, like IOTA’s Tangle [56, 57] based on Directed Acyclic Graphs. However, there are remaining challenges for achieving desired privacy, security, and efficiency goals [58, 59].

To properly fulfill the privacy and security goals in SSI, PETs should be applied [34]. There exist many cryptographic primitives that help increase privacy, like group [60] and ring [61] signatures, pseudonyms [11] or zero-knowledge proofs [62]. From these, multiple PETs have been developed. For instance, Direct Anonymous Attestation (DAA) [63] enables authentication of messages through signatures, while preserving privacy in a pseudonymous way (controlled linkability). Another example, now considering digital identity as a collection of attributes, is Attribute-Based Signatures (ABS) [64]. ABS allow the signing of a message according to the signer’s attributes. That is, verification of the signature will only be successful if the signer’s attributes fulfil a specific attribute-based policy.

A similar notion is privacy-preserving Attribute-Based Credentials [10, 11], which is the central focus of this thesis due to its alignment with privacy-preserving identity management, particularly within the context of SSI principles. Current initiatives like the EUDIW [8] note this synergy, even if current implementations do not include p-ABCs. In essence, p-ABCs enable the issuance of credentials to users, who can then derive unlinkable one-time tokens from them, revealing only part of the information. There exist identity management systems that rely on p-ABCs for privacy-preserving authentication, such as

---

<sup>1</sup><https://data-spaces-business-alliance.eu/>

<sup>2</sup><https://hub.ebsi.eu/>

uProve [65] or Identity Mixer [66]. The latter has been used in various European projects like ABC4Trust [29] and ARIES [30], and some solutions like IRMA authentication [67]. Similarly, “ePASSO” [68] is a SSO-like authentication solution that relies on p-ABCs and improves usability while ensuring security and privacy. These solutions offer better user privacy than traditional SSO systems, but they still rely on a single IdP that becomes a single point of failure. Additionally, their real-world impact has been very limited, partially due to efficiency, usability, interoperability, or lack of advanced security features that justify the additional burden of p-ABCs, as will be expanded in Section 2.1.2.

All in all, digital identity management has been a key challenge since the advent of online interactions. Traditional systems and particularly SSO have become widespread, but present glaring privacy issues. Solutions applying SSI are not yet mature but their importance is steadily rising. To reach meaningful privacy notions, it will be important to rely on PETs like p-ABCs. However, existing systems have not covered the necessary features in efficient, secure or usable ways, hampering their success. The solutions discussed in this thesis aim to address these challenges by providing extensible, user-friendly implementations of p-ABC systems that take into account the current landscape of identity for ensuring interoperability and applicability to real-world scenarios.

### 2.1.2 Privacy-preserving attribute-based credentials

As previously introduced, p-ABCs [10, 11], also known as anonymous credentials, take up our focus as they pose an attractive building block for privacy-preserving identity management. Indeed, they provide significant security and privacy guarantees, along with advanced features through zero-knowledge proofs tailored to attribute-based authentication use cases.

Informally, p-ABCs are generally required to fulfil two security properties: *unforgeability* and *privacy* (or *unlinkability*). Here, *unforgeability* refers to the inability of adversaries to present attributes without access to a valid credential that contains them. Whereas *privacy* refers to the inability of an adversary to obtain information from a presented token apart from what was intended to reveal by the user. Usually, this notion includes the impracticability of linking multiple presented tokens generated from the same credential (i.e., *unlinkability*). However, the *privacy* requirement can be relaxed to ensure that tokens cannot be linked to a specific credential, but multiple showings will be linkable between themselves [69]. This definition covers single-show credentials such as [70], in contrast to multi-show credentials like [28] that are usually more interesting for general purpose identity management.

#### Single issuer and multi-issuer credentials

Since their inception, numerous general-purpose p-ABC schemes have been proposed, many of which are based on Camenisch-Lysyanskaya (CL) [71, 72] signatures. They have been instantiated in multiple solutions, such as IBM’s identity mixer (Idemix) [28, 29, 73] or Microsoft’s UProve [65, 70]. In recent years, various works have focused on other signatures schemes such as Pointcheval-Sanders (PS) [74, 75] and BBS [76, 77], mainly

because of their efficiency (e.g. [78–83]). The majority of existing p-ABC solutions rely on a single credential issuer, introducing similar single-point-of-failure vulnerabilities as in traditional SSO systems. Several cryptographic solutions can be employed to mitigate the issue. For instance, blind signatures [72, 84, 85] can be used to hide attribute information from the issuer. However, in practice, this often requires a complex and computationally heavy issuance processes with zero-knowledge proofs [86] to ensure the validity of the signed information. Additionally, the single point of failure issue is not fully addressed. Therefore, distributed cryptographic tools such as distributed oblivious pseudorandom functions [87] or multi-party computation frameworks [88] may be applied to tackle this problem. For instance, this has been explored in the context of SSO systems [87, 89].

In this regard, solutions like Coconut [90], for PS signatures, or Doerner *et al.* [91], for BBS+ signatures, enable threshold issuance. In these schemes, a subset of  $t$  out of  $n$  issuers is needed to generate a credential. However, this setup introduces the challenge of a trusted or complex process for key generation, with difficult (or unrealised) ways of dynamically changing keys or signers. A similar scheme based on CL signatures that fixes  $t$  to  $n$  so that every participant is needed for obtaining a signature is introduced in [92]. Meanwhile Camenisch *et al.* [15] introduce a variant of the PS p-ABC scheme based on multi-signatures [93]. Multi-signatures are a particular type of distributed signatures where each party can independently sign messages, allowing for the compression of signatures generated by multiple parties on the same message into a single compact signature. An equivalent notion is achieved by Héban and Pointcheval [78] with their proposal for multi-authority credentials.

### Beyond selective disclosure

While selective disclosure and unlinkability typically stand out as key goals when applying p-ABCs, these schemes are versatile and may offer advanced security and privacy features or efficiency optimisations tailored to the specific requirements of the identity management use cases. Since their initial introduction, a wide range of extended functionalities have been suggested. In the following, we give a brief overview on some of the most prominent ones. Revocation [94, 95] enables the invalidation of a certificate. Inspection [78, 96, 97] allows a trusted party to re-identify the user (or credential) that computed a presentation token. These two features serve as critical security measures against credential misuse or theft. Range proofs [98–101] allow proving predicates over numeric attributes without revealing the precise value, with the archetypal example being proofs of age. Set-membership proofs [102, 103] serve a similar purpose, but focus on proving that the attribute is within a set without revealing the actual value. Pseudonyms [69] give the user the opportunity to control the linkability of the generated tokens, e.g. enabling repeated authentication against the same party while remaining unlinkable in any other context. Moreover, there has been interest in schemes that are tailored to specific scenarios. For example, [104] proposes *updatable* credentials, with a direct application to incentive systems. *Delegatable* credentials [105–107] create a hierarchy of trust where credential holders can delegate responsibilities to other participants. *Issuer-hiding* credentials [108] enable the verification of the presentation token without revealing which issuer within a group generated the credential, avoiding indirect leakage of user attributes implicitly learned

through the issuer (e.g., nationality). Habock *et al.* [80, 109] focus on delegating the bulk of the management of p-ABCs to external services.

While the theoretical work on such features has been extensive, there has been a significant lack of implementation efforts supporting them in an extensible, efficient and usable manner. Works like [15, 78, 81, 90, 91] have reduced implementations that miss support of operations beyond selective disclosure. Broader implementation projects like Hyperledger’s Aries<sup>3</sup> and the related cryptographic library Ursa<sup>4</sup> also focus solely on basic protocols with selective disclosure, considering attribute predicates or distributed issuance as out of scope. IRMA authentication [67] is based on Idemix, but concentrates only on selective disclosure and revocation tied to a specific application, and cannot be used as a re-usable p-ABC software.<sup>5</sup> Likewise, the implementation<sup>6</sup> of the anonymous credentials presented in [110] offers pseudonyms (generated by the issuer) and revocation functionality. The implementation<sup>7</sup> used in the H2020 project ABC4Trust [111] is designed for the Idemix library. The implementation supports various predicates in a modular way, including pseudonyms, revocation or inspection. However, its efficiency is notably lacking, the code base has been abandoned, and it does not consider the distributed issuance scenario. A recent paper by Rosenberg *et al.* [112] uses general purpose ZK-SNARKs to implement p-ABC functionality. This ensures the flexibility to include any type of predicates (that can be computed efficiently) over attributes. However, it focuses on using bulletin boards to publish credentials and information needed for their proving and verification, disregarding practical considerations such as communication time. While its use of general-purpose ZK-SNARKs can also function with traditional ABCs, this introduces a higher computational burden on provers and additional practical burdens like the generations of the common reference string.

### Non-transferability

On a different note, one of the main drawbacks of many existing p-ABC systems, particularly those that are solely software-based, is the possibility of copying, sharing or selling credentials. Thus, a challenge of special interest for their adoption in real-world systems is achieving privacy-preserving *non-transferability*.

The literature presents various approaches to address this issue, each with their own strengths and limitations. Camenisch and Lysyanskaya [97] proposed *all-or-nothing sharing* as a deterrent against credential sharing. In their scheme, lending a credential to another entity leads to the credential being available for them in any context, effectively assuming control over the user’s identity. However, losing control over a credential may have a limited impact on the legitimate user depending on its contents (e.g., a “Green Pass” certificate<sup>8</sup>) or their level of trust in the other party (e.g., close family or friends).

<sup>3</sup><https://www.hyperledger.org/use/aries>

<sup>4</sup><https://www.hyperledger.org/use/ursa>

<sup>5</sup><https://github.com/privacybydesign>

<sup>6</sup><https://github.com/hyperdivision/anonymous-credentials>

<sup>7</sup><https://github.com/p2abcengine/p2abcengine>

<sup>8</sup>[https://ec.europa.eu/info/live-work-travel-eu/coronavirus-response/safe-covid-19-vaccines-europeans/eu-digital-covid-certificate\\_en](https://ec.europa.eu/info/live-work-travel-eu/coronavirus-response/safe-covid-19-vaccines-europeans/eu-digital-covid-certificate_en)

Another approach involves relying on tamper-proof hardware to protect the usage of p-ABCs [113, 114], akin to the TPM binding used in DAA [63, 115]. The credential will only be usable on the adequate device, preventing its duplication. However, this mandates the use of dedicated hardware for each user, or additional management such as re-issuing credentials after a device change if “general-purpose” devices like mobile phones are used. This reduces their flexibility and usability in various scenarios. Further, sharing the device itself, e.g. with close relatives, may still be an option for users leading to credential misuse.

A solution to the issue of credential lending is to bind the physical identity of the user to the credential through biometrics. In this way, only a user that proves she has matching biometrics during the presentation phase will be able to use the credential in an authentication process. This approach has been realised with different techniques in the literature, and is particularly applicable to privacy-preserving physical access control. For instance, [116] proposed solutions based on a dedicated trusted device, e.g., a smart card, carried by every user. The device is trusted by all parties to scan a fresh biometric (e.g., a fingerprint) of the user and verify that it matches the stored one which was bound to the credential issuance. Blanton and Hudelson [117] propose an equivalent approach but the biometrics are handled through Fuzzy Extractors (FE) [118, 119] to avoid storing biometric readings. On a high level, FEs take as input a sample from a noisy source (e.g., biometric data), and output the same digest as long as the two samples are sufficiently close to each other. Later works like [120, 121] have taken advantage of the deterministic output of FE to avoid explicit matching in the trusted device. Instead, the trusted device produces outputs based on biometric sampling later treated through a fuzzy extractor, which are encoded into the credential. The sensitive results never leave the trusted device, and are in fact deleted after each process. Thus, a user needs to interact with the trusted device during each presentation process, which will only generate the necessary outputs if the biometric reading is correct.

Nevertheless, these solutions suffer from the same usability limitations as device-bound credentials, with increased scalability issues because they require dedicated hardware per user. Additionally, while FEs are an attractive object [118, 122–125], their practical application is hindered by the loss of entropy that leads to accuracy levels ( $\ll 90\%$ ) significantly below contemporary biometric systems [126, 127]. Moreover, their auxiliary data requires significant storage (or bandwidth when transmitted), often reaching hundreds of megabytes or even gigabytes [125]. Hesse *at al.* [128] introduce the concept of credentials with visual holder authentication. The setting integrates an additional physical device in the credential showing process, which is capable of displaying a picture of the holder for manual verification by the verifier. Finally, Adams [129] explored a different angle based on one-show credentials [70], where the biometric sensor encrypts the measured biometrics for the user, and hands a commitment to the value to the verifier. The user then computes a zero-knowledge proof of knowledge showing that the biometrics certified in the credential match those in the commitments. However, security is argued in an ad-hoc manner without formal treatment and Adams concluded that his solution is inefficient for practical applications.

Altogether, p-ABCs have received considerable focus in the literature for privacy-

preserving applications. However, their uptake in practical implementations and real-world applications has been severely limited. This is particularly noticeable when considering advanced features that make p-ABCs appealing despite their increased complexity over traditional cryptography, such as distributed issuance to avoid the single-point of failure issue prevalent in identity management solutions, or extensible application of security and privacy predicates like revocation or pseudonyms. Moreover, approaches to address *non-transferability* have lacked formal treatment, disregarded practical aspects for their implementation like suitable biometric features or performance evaluation, or had limited application scope because of their characteristics (*there is no silver bullet approach*). The solutions explored in this thesis aim for practical p-ABC systems with distributed issuance, modular and extensible realisation of advanced features like inspection or range proofs, and real-world applicability tackling key past issues such as interoperability or integration into an identity and trust framework. Furthermore, we formalise a biometric approach for non-transferable p-ABCs and show its practical value for risk-based access control (RiBAC) [130].

### 2.1.3 IoT identity management

The challenge of identity management also extends to the realm of the IoT. In the IoT field, unique obstacles arise due to factors such as vast numbers of devices, computational constraints or heterogeneity. Traditional approaches relying on Public Key Infrastructure (PKI) and X.509 certificates are in a mature state, with many commercial solutions like [131]. However, concerns of scalability, complexity and usability of such systems are considerable in IoT environments, particularly with their increasing prevalence even in environments with non-expert users. This jeopardises the applicability of such centralised solutions, although the state of art presents directions to address those issues, such as lightweight designs of PKI protocols [132]. Nonetheless, current trends follow the general tendency in identity management that veers from traditional PKIs towards solutions that better cover privacy and usability goals.

One such initiative is the FIDO's Device Onboard Specification (FDO) [133], which involves multiple device manufacturing companies. This work focuses on the initial phases of the device's lifecycle. It tackles *bootstrapping* of the device identity after manufacturing, and performing *secure bootstrapping in the target domain*. Similarly, contributions like [134, 135] aim to improve the security of those phases by providing *roots of trust* for identifying a device based on Physical Unclonable Functions (PUF).

The literature also reflects a growing interest in applying SSI principles and technologies to IoT scenarios. For instance, Kortensniemi *et al.* [136] assess the viability of using DIDs as an *identification method within a domain*. Various studies propose adaptations of this approach [137–139], disregarding IoT challenges like initial bootstrapping. Diego *et al.* [140] present some concepts on initial provisioning and consider Attribute-Based claims focusing on a specific application scenario. This is similar to the work in [141], where authors focus on identity in a highly controlled environment. In summary, multiple studies apply SSI technologies but neglect to address crucial privacy principles such as complete control over identity or ensuring privacy throughout authentication processes.

Various works study practical aspects of the application of decentralised technologies, particularly DLTs such as blockchain, to the IoT field as enablers for identity management [57, 142, 143]

Other research in the literature addresses privacy concerns in a more direct manner. Gu *et al.* [144] present a methodology based on mesh signatures [145] that allows devices to participate in a domain while preserving their anonymity during interactions. Their approach relies on a centralised architecture and neglects cases where identification of specific devices is required. Yang *et al.* [146] achieve a similar objective leveraging accumulators [147], outsourcing witness update computations enhance efficiency. Choudhury [148] proposes a scheme for authenticating in 5G mobile networks in an efficient, secure, and privacy-preserving manner. The solution relies on a federation toward the home 5G network. This renders the approach as less apt for general privacy-preserving authentication, but poses appealing security and privacy properties for specific cases such as bootstrapping in mobile networks.

In the discussed instances, as well as in other works in the broader landscape of IoT identity management such as [149], the emphasis is put on the issue of (pseudonymous) identification. A crucial element for privacy-preserving identity management absent in these works is the facilitation of *attribute-based authentication/authorisation processes*. As argued above, p-ABCs are a particularly useful tool to achieve this goal with privacy and security guarantees. Their application in IoT contexts has been even more limited than in general-purpose applications. This limitation stems mostly from the higher relative impact of p-ABCs computational demands and lack of interoperability with traditional solutions. Nevertheless, a few works in the literature have explored the usage of p-ABCs to achieve privacy goals in IoT identity management. For instance, Canovas *et al.* [150] advocate for using Idemix [28, 29, 73] in IoT devices for attribute-based authentication. To mitigate its inefficiency, they delegate some of the computational work to an external server. The parallel work by the same authors [151] additionally considers the security notion of inspection. Despite these efforts, the inefficiency of the solution even with delegated computations makes it impractical for most scenarios. Moreover, it primarily focuses on technical aspects of the application of p-ABCs, neglecting its relationship to other IoT identity management challenges like the device lifecycle. Alcaide *et al.* [152] also fall short in addressing the dynamic nature of IoT environments and lifecycle management. Further, [153] identifies a successful attack on their approach. Lastly, Neisse *et al.* [154] also explore an Idemix-based approach in an IoT context, yet their evaluation focuses on user machines rather than IoT devices themselves.

In definitive, the literature shows the high interest of the research community in achieving highly automated, secure and privacy-preserving identity management in IoT scenarios. There are various works tackling different challenges attached to this objective such as providing roots of trust for the identity, enabling identification of devices or authentication and authorisation processes based on identity attributes. However, the existing solutions often fail to properly address key privacy goals. More importantly, they neglect to achieve meaningful links and coherence between solutions to these challenges, which must all be addressed for achieving real-world applicability. The solution developed within this thesis focuses on achieving authentication and authorisation processes

aligned with SSI principles, while considering their link to the initial bootstrapping and enrolment phases for a holistic view on the identity management lifecycle. It also gives insights on several practical considerations of such identity framework and its application to a validation scenario.

### 2.1.4 Gap analysis

As hinted in earlier sections, the analysis of the literature related to the diverse studied challenges lead to the identification of various gaps in existing research. Traditional identity management systems like SSO schemes [38, 39] have critical flaws in terms of privacy, particularly in terms of user tracking and adhering to principles like minimal disclosure. While the rise of SSI [7] results in a better landscape for users, current solutions have not fulfilled the security and privacy requirements, largely because of a lack of implementation of adequate privacy-enhancing technologies. The few existing privacy-preserving identity management systems relying on p-ABCs have not addressed the single-point of failure issue for the identity provider. Moreover, their impact has been limited due to poor efficiency, usability, and lack of interoperability with identity ecosystems [65–68]. Thus, we identified a gap for a privacy-preserving identity management system that avoids a single point of failure, provides usable and extensible access to p-ABCs, and is inline with SSI security and privacy principles and technologies such as DLTs as a trusted verifiable registry.

In the field of p-ABCs, while some reduced implementation for schemes with threshold issuance exist [90, 91], the theoretical groundwork of fully distributed p-ABCs (e.g., Pointcheval Sanders Multi-Signatures [15]) has not lead to a full-fledged implementation, nor an integration and evaluation in a practical system. Moreover, advanced features that differentiate p-ABCs from simpler solutions and justify their increased complexity have not been taken advantage of. While the achievement of these advanced features in a theoretical sense has been thoroughly studied [78, 94–101], practical implementation has been very reduced, and particularly never paired with distributed issuance [66, 110–112]. More critically, no implementation of identity management with p-ABCs offers efficient instantiation of such functionalities in a modular and extensible way. This would enable the introduction of advanced privacy features into heterogeneous ecosystems, with a clear way to adapt and improve them whenever necessary. Lastly, the issue of non-transferability has lead to various approaches like binding credentials to secure hardware or biometrics, but it has become apparent that there is no silver bullet approach. The latter is particularly appealing for physical access control for use cases like risk-based access control, e.g. during pandemics [130]. However, most existing solutions rely on deploying trusted hardware per user, which does not fit many scenarios especially in terms of scalability. There are no practical instantiations that avoid this issue. What is more, the few works that have considered this approach lacked a formalised security framework [113, 114, 116, 117, 121, 128, 129].

Finally, the analysis of the literature shows that researchers are aware of various challenges in IoT identity management. Providing a root of trust for device identities, enrolling devices in domains, or enabling operational flows with secure and privacy-preserving iden-



tification and authorisation are necessary for solutions that protect the whole lifecycle of the device. Existing research handle one or a few of these challenges, with varying success in terms of privacy or security, but none of them cover the complete set [133–141, 144–146, 149–152]. Therefore, there is lack of a framework enabling automated, flexible, secure and privacy-preserving identity management that holistically covers challenges throughout devices’ lifecycle in a coherent and interlinked manner for heterogeneous IoT scenarios.

Summarizing, we identified the following gaps:

- G-1** A privacy-preserving identity management system that provides usable p-ABCs and avoids the single point of failure.
- G-2** A dp-ABC system that offers advanced features such as inspection or pseudonyms in a modular, extensible and efficient way.
- G-3** A comprehensive decentralised zero-trust framework that covers identity following SSI principles and offering privacy capabilities through p-ABCs.
- G-4** An instantiation of p-ABCs taking advantage of their privacy capabilities in an interoperable way with widespread specifications, particularly Verifiable Credentials.
- G-5** A formalisation and instantiation of practical non-transferable p-ABCs based on binding biometrics to credentials that avoid dedicated hardware per-user.
- G-6** A practical application of p-ABCs to enable privacy-preserving authorisation scenarios in IoT.
- G-7** A framework that comprehensively addresses privacy-preserving identity management challenges throughout IoT devices’ lifecycle.

## 2.2 Results

This section summarises the main results obtained through the development of the thesis with the aim of covering the identified gaps and achieving the established objectives. In the following, we give an overview of these results, while following sections delve in deeper detail into the publications of the compendium.

Initially, we developed an implementation of distributed privacy-preserving Attribute-Based Credentials (dp-ABC) based on Pointcheval-Sanders Multi-signatures (PS-MS). This implementation was empirically evaluated in the article [18], of which further details can be found in Section 2.2.1. This evaluation focused on measuring the execution time of the scheme’s methods and assessing their behaviour and scalability in terms of the relevant parameters, showing a significantly superior performance over Idemix. Additionally, this work described how the implementation was used to enhance a distributed, practical and user-friendly identity management system within the framework of the EU H2020 project OLYMPUS addressing G-1. This system, designed for SSO-like usability, defines a transparent “virtual identity provider” comprised of partial identity providers that enables user-centric privacy by design in delegated identity, as described in [17, 25],

addressing G-1. The identity management framework was later expanded to achieve a privacy-preserving zero-trust access control framework, whose application to a smart city platform designed with user privacy and control at its core can be found in [22].

Advancing the distributed nature of the solution, and aligning with emerging SSI trends, the identity management framework was enhanced with distributed ledger technologies to foster a decentralised trust ecosystem [155]. Along with the previous outcome, this tackles G-3. The main idea is moving away from traditional centralised PKIs for the dissemination of public keys and other data necessary for interacting with the distributed identity provider. This was realised through a blockchain-based deployment, relying on smart contracts to register and resolve data in an automated way. Particularly, we define smart contracts that automatically computes key aggregation for distributed identity providers comprised of multiple partial issuers. Additionally, the registered data will include information about each issuer, such as their individual public keys, and public parameters like the definitions for attributes in issued credentials. Due to blockchain properties, this creates an immutable record of who are the identity providers in the ecosystem, their characteristics, and how they were created, acting as a source of trustworthiness in the framework. Furthermore, the ecosystem allows for building trust relationship among its participants, broadening the scope to service providers. Service providers willing to participate in the ecosystem will be discoverable through the blockchain registry. For this goal, they will be required to announce the policies along with services offered, including the rationale and purpose. This produces an auditable registry of services, enabling the analysis of their behaviour. For instance, users or monitoring entities may be able to notify malpractices such as excessive request of data in discordance with currently proclaimed policies, or the evolution of service offerings and policies over time. Thus, the trustworthiness participants can be continuously examined and asserted, leading to a safer ecosystem.

Focusing on enhancements to the dp-ABC scheme, the work outlined in [21] develops a refined zero-knowledge presentation protocol that enables modular extensions through commit-and-prove techniques, addressing G-2. This outcome is validated by implementing and evaluating four advanced features commonly sought in p-ABCs: range proofs, pseudonyms, inspection and revocation. The results of the evaluation show the practicality of the approach, whose suitability for improving the current identity management landscape is discussed. We defer more details to Section 2.2.2. The practical application and integration of this solution into the aforementioned identity management system were further validated in relevant use cases within the context of the EU H2020 Cyber-Sec4Europe project [26, 27].

The challenge of achieving *non-transferability*, another prevalent issue for practical p-ABCs, was also tackled during the development of the thesis. We focused on physical access control scenarios where private, secure and scalable verification is required, such as in Risk-based Access Control during pandemics. In this context, we introduced [23] a formalised framework for Biometric-Bound Attribute-Based Credentials (bb-ABC). In this approach, instead of relying on a trusted dedicated device per user, each verifier is equipped with a semi-trusted biometric reader. Within this framework, we developed and detailed two constructions, proving them secure. Furthermore, we described practical in-

stantiations of these constructions, showing their efficiency and suitability for real-world applications. More details on these achievements and how they tackle G-5 are provided in Section 2.2.3.

During the thesis, significant effort has been devoted to mitigating a key limitation of existing p-ABC schemes: their lack of interoperability with other solutions and emerging trends in identity (cf. G-4). This challenge has been part of the focus of the OLYMPUS identity management system, with the integration of dp-ABCs in an SSO-like identity provider transparently for users [17, 18]. With this objective in mind, we developed an integration of p-ABCs and the W3C’s VC specification, complemented with a blockchain implementation acting as a verifiable data registry for the necessary public information through smart contracts. The aim was two-fold: make the implemented dp-ABC solution compatible with the standard, and thus transparently integrable into solutions based on it; and provide a reference integration that tackles practical issues of applying p-ABCs to the specification, paving the way for privacy-preserving alternatives to the existing VC instantiations. To achieve this, we define a set of materials and guidelines for the integration of p-ABCs with the specification. This includes JSON-LD contexts for defining appropriate fields and (meta-)schemas for facilitating the creation and validation of credential schemas containing the necessary information for zero-knowledge proofs, e.g., by establishing types for attributes, or minimum and maximum values for numeric attributes. These efforts culminated in the definition and implementation of a *signature suite* for the PS-MS scheme, introducing two new proof types for the scheme’s signatures and zero-knowledge tokens. In this way, the proofs can be generated and verified transparently along with other signature suites like *Ed25519* signatures. This work has been published through several contributions [19, 20, 155], and more details can be found in Section 2.2.4 along with the results of [19].

These advancements on interoperability have been key for achieving the final outcome of the thesis, which was published in [19]. In this work, we define, analyse and instantiate a comprehensive privacy-preserving identity management framework for IoT environments, covering the whole lifecycle of devices (c.f G-7). The dp-ABC solution is proposed as an enabler for authentication and authorisation processes, showcasing its efficiency. Its interoperability is fundamental to ensure the flexibility necessary for accommodating the heterogeneous participants in the IoT ecosystem, from services and end-users to resource-constrained devices. Thus, it is apt to address G-6. The framework also addresses challenges like establishing a root of trust for device identities or bootstrapping and enrolling in a security domain, establishing coherent links between the steps. This publication, further delineated in Section 2.2.4, also describes an instantiation of the framework in the context of the EU H2020 ERATOSTHENES project. Beyond identity management, the project architecture encompasses trust and security lifecycle management of devices within security domains. The instantiated identity solution is a cornerstone for the architecture and, conversely, it is enriched by its integration with the other technologies. For instance, trust management complements the adaptability of the identity framework by enabling adapted responses to diverse alternatives for identity mechanisms, such as establishing increased trust scores for devices with *stronger* roots of trusts like physical unclonable functions [24].

### 2.2.1 Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures

The implementation and evaluation of the distributed p-ABC (dp-ABC) scheme based on PS-MS [15] introduced in [18] served as a crucial building block for the research developed in this thesis. The scheme, which had not seen any practical instantiation or validation yet, can be informally described as a set of eight algorithms grouped in three phases. First, the **setup** phase is comprised by a *setup* method for establishing public parameters like underlying groups, as well as the methods for *generating key pairs* for signing and *aggregating public keys* into a single verification key associated to a group of signers. The next phase corresponds to credential **issuance**, and it has the usual *signing* and *verification* processes. Additionally, the *combine* method allows the aggregation of signatures over the same set of attributes generated by different issuers. Lastly, during the **presentation** phase, the user can *generate zero-knowledge proofs* to disclose a subset of attributes. Any entity will be able to *verify the zero-knowledge proof* against the issuer's (whether a single or a group) public key.

The aforementioned aggregation of keys and signatures is one of the key features of the scheme, as contrary to existing dp-ABC implementations both of these methods can be carried out without any interaction between signers. This allows avoiding the single-point of failure of the issuer without having to rely on complex setup processes or trust relationships between issuers. In fact, this property can be applied to enhance decentralised use cases by stretching the boundaries of the zero-trust [156] paradigm to encompass identity management itself. That is, users will possess credentials certified by different issuers, and the trust deposited on their identity will depend on the aggregated trustworthiness of the set of issuers. The caveat is less resilience against denial of service than solutions based on threshold issuance.

The implemented scheme was integrated into the solution developed in the H2020 OLYMPUS project, which devises a privacy-preserving identity management solution employing distributed cryptographic techniques. Here, the traditional identity provider is substituted for a virtual identity provider (vIdP) composed of  $n$  individual IdPs. The OLYMPUS architecture [17] integrates user authentication through distributed password verification, aligned with traditional SSO solutions for ensuring usability (Figure 2.1 shows the simplified architecture of the solution). Additionally, it integrates two solutions for authentication against relying parties. First, distributed token generation based on the PESTO scheme [157], which completes a distributed SSO solution. Second, p-ABCs for offline authentication, issued in a distributed manner. The requirements coming from this architecture shaped the goals of the implementation, with extensibility or configurability (e.g., of security level through the underlying curve) being prioritised. The seamless integration into the system provides a user-friendly way to use dp-ABCs in practical applications.

The implementation was thoroughly evaluated through empirical experiments. First, the efficiency of each method along with its behaviour and scalability with respect to the relevant parameters was measured. For instance, Figure 2.2 shows measurements for experiments in the issuance and presentation phase, which are the most common processes

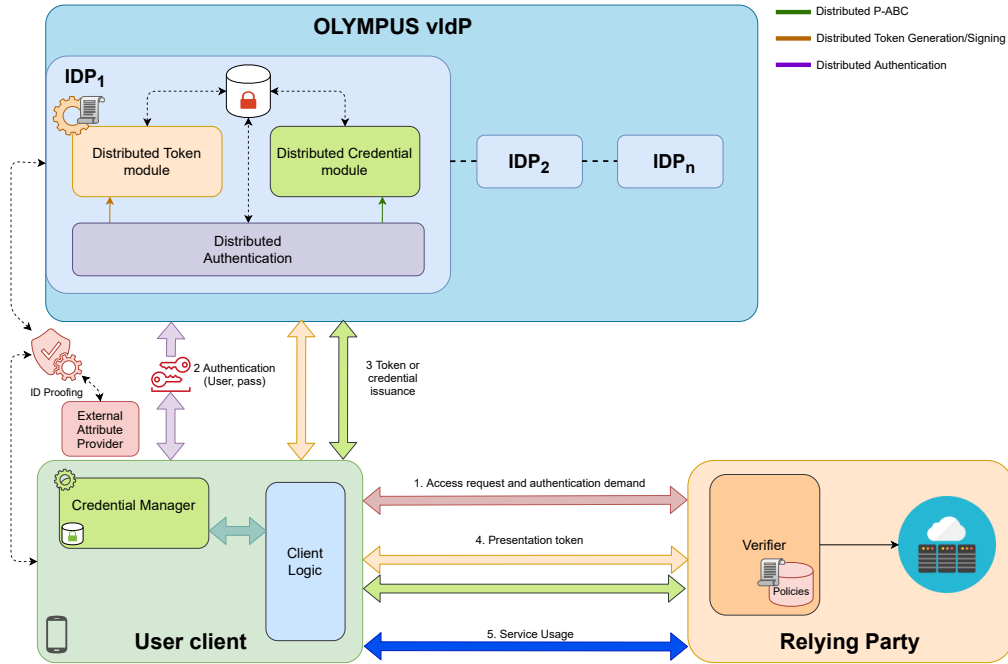


Figure 2.1: Simplified version of the architecture of the OLYMPUS project [17]

involving users. The obtained values denote the practicality of the implementation. The most expensive method, and the only one that can grow super-linearly when multiple parameters are modified (attributes and number of signers), is key aggregation. Nonetheless, it is usually a one-off operation, and still 300 milliseconds in the worst case. The issuance process, including signing, signature aggregation and verification, grows linearly with the number of attributes and signers. The presentation phase grows linearly with the number of attributes as well, with the caveat that proof generation time only depends on the number of hidden attributes. In all cases, both issuance and presentation take significantly less than 100 milliseconds in total.

These encouraging results correspond to the instantiation of the scheme with BLS12-461 as the underlying pairing-friendly curve. This curve is estimated to provide 134 bits of security [158], which is desirable for very strict security requirements. If the bar is slightly lowered to just under 128 bits [158], using the BLS12-381 curve leads to a 40% decrease in execution time. A similar improvement can be achieved again by using the BN254 curve, but this would mean reducing the security guarantees to  $\sim 100$  bits after a recent attack was discovered [159], which should rule out their use in practice.

Lastly, the implementation was compared with Idemix [29, 30, 73], as shown in Figure 2.3. The comparison only involved equivalent functionalities of both implementations. Thus, the complete issuance processes of the schemes were directly compared, including distributed issuance for the PS-MS scheme. However, attribute blinding was disregarded in the Idemix method as it is not considered in our implementation. Additionally, only selective disclosure proofs were considered for the presentation phase.

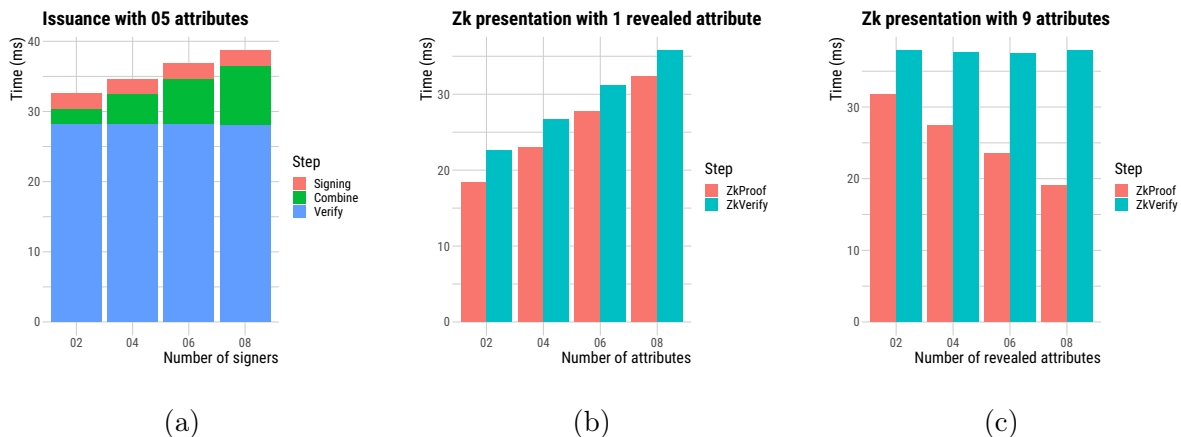


Figure 2.2: Execution time of PS-MS with BLS12-461 for (a) Issuance phase, (b) ZK presentation (with respect to the total number of attributes) and (c) ZK presentation (with respect to the number of revealed attributes)

The comparison shows a significant discrepancy between the mean execution times, in terms of base values (between twice and four times slower) and rate of growth. In light of significant variances for the execution times in Idemix, the statistical significance of the results was corroborated through non-parametric permutation tests, obtaining  $p$ -values below the  $10^{-4}$  mark for all parameters. What is more, these values were obtained with an Idemix configuration that achieves 112 bits of security [160], while the PS-MS scheme had the underlying BLS-461 curve, reaching 134 bits.

## 2.2.2 Beyond Selective Disclosure: Extending Distributed p-ABC Implementations by Commit-and-Prove Techniques

An extension to the dp-ABC implementation was proposed in [21]. Particularly, we developed a formalised modification of the presentation phase, which allows for modular extensions of the token generation through commit-and-prove zero-knowledge proofs. The main idea to achieve this is splitting the proof goal into the individual goals for each predicate along with a proof goal for the showing of the credential itself. To do so, we modify the presentation protocol of the scheme to enable "linking" attributes to Pedersen commitments. We prove that this modification of the scheme maintains the same security properties as the original. That is, the method is still a non-interactive zero-knowledge proof of knowledge. The security of the overall decomposition is then simple to verify, as soundness and zero-knowledge can be reduced to the equivalent properties of the sub-proofs along with the binding and hiding properties of Pedersen commitments (cf., e.g., Benarroch *et al.* [102] for formal details). Nonetheless, we additionally establish that all sub-proofs will be bound to a *context* that includes an unambiguous serialisation of the proof goal. With this, we ensure the non-malleability of the proof goal, avoiding any potential mix-and-match attack.

Through this modification, we build an improved dp-ABC system with advanced pri-

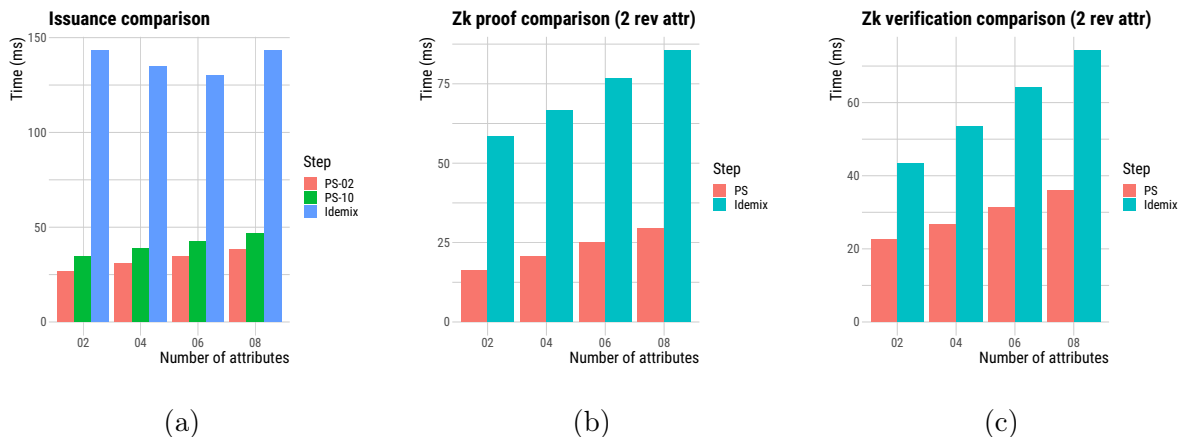


Figure 2.3: Execution time comparison between Idemix and PS-MS with BLS12-461 for (a) Issuance (2 and 10 signers for PS-MS), (b) ZK proof generation and (c) ZK proof verification

vacy and security features, as shown in Figure 2.4. Particularly, we implemented and evaluated the practicality of the scheme enhanced with four applications: range proofs, pseudonyms, inspection and revocation. As depicted in the flow, a new formal entity is introduced in the system for governing each of the latter two. We give further insights on these functionalities below.

Through **range proofs**, the system supports fine-grained predicates for numeric attributes, contributing to the goal of minimal disclosure. Namely, attributes can be proven to lie in a range without revealing the specific value. This enables, e.g., commonplace policies such as age proofs or non-expiration of a credential. Among the various approaches in the literature, we opted for Bulletproofs [98] to instantiate these proofs because of their relative efficiency and small proof size (cf. Figure 2.5a). Furthermore, the implicit constraint of credential attributes to relatively small ranges of valid values (with respect to integer modules used in secure settings) allows for using Bulletproofs for arbitrary ranges through *offsets* [161] without needing two executions of the protocol in one-side inequalities.

The introduction of scope-exclusive **pseudonyms** gives users fine-grained control over their linkability [69]. Users may re-authenticate to a service using the same pseudonym, while keeping unlinkability across providers. The scopes can also be time-gated, allowing use cases such as privacy-preserving detection of account sharing. In order to achieve the functionality, we apply the common approach of embedding in the credential an attribute that is never disclosed, but is used along with the scope to derive the pseudonym instead. For computing the pseudonym, we implement the same approach as [69], through hashing the scope into the adequate group and using the secret attribute *id* as an exponent, as shown in Figure 2.5c.

**Inspection** offers a trade-off between full anonymity and accountability when using p-ABCs. A trusted entity may re-identify users (e.g., through recovering an identifier) from a transcript of the presentation, for instance in case they breach agreed policies.

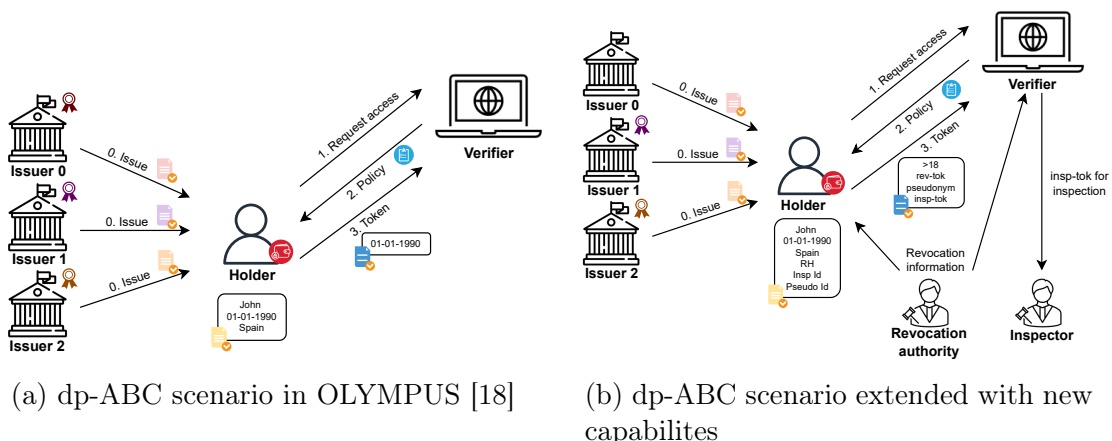


Figure 2.4: Scenario and flows supported by new dp-ABC functionality [21] (b), in contrast to previous solution [18] (a)

Following the usual methodology, we instantiate the functionality through encryption of an *id* attribute under the public key of the inspector (cf. Figure 2.5b). The user will prove upon presentation that the encrypted value corresponds to the attribute embedded in the credential. Note that, thanks to the modularity of the solution, the inspection process relies exclusively on the specific sub-proof. Moreover, in line with the goals of the dp-ABC scheme, this process can be carried out in a distributed manner through multi-party computation of the decryption [88, 162].

Lastly, **revocation** offers the possibility to invalidate a credential (e.g., identified through inspection) after theft or misuse. To achieve this notion, we implemented an allow-list approach, with the revocation authority maintaining an implicit list of the credentials valid in a specific *epoch*. The authority realises this by issuing its own credentials on a secret attribute *rh* that is never revealed during presentation. To avoid losing unlinkability, we rely on the same p-ABC implementation for the revocation credentials (cf. Figure 2.5d). This approach is suitable for scenarios with low occurrences of revocation, and avoids the high computational complexity for users in accumulator-based solutions [163, 164]. Moreover, as in inspection, distributed techniques can be applied to avoid the single-point of failure of the revocation authority.

The overhead of these functionalities over the basic scheme was measured for the proving and verification processes, including cryptographic operations and related processes such as policy parsing. The results are summarised in Figure 2.6. The overhead incurred by one inspection or pseudonym predicate is less than 20 milliseconds for both proving and verifying, while revocation adds around 30 milliseconds to each operation. These values are valid for practical applications. However, range proofs are much more expensive, especially on the prover side. Nonetheless, proofs for numeric values that can be represented in fewer than 32 bits, which cover most use cases, take less than 1 second in total. Even in the case of 64 bits, the total execution time falls under 2 seconds. We remark that, while Bulletproofs are widely used because of their overall efficiency, other proofs with different characteristics like lighter costs on prover side might be useful for



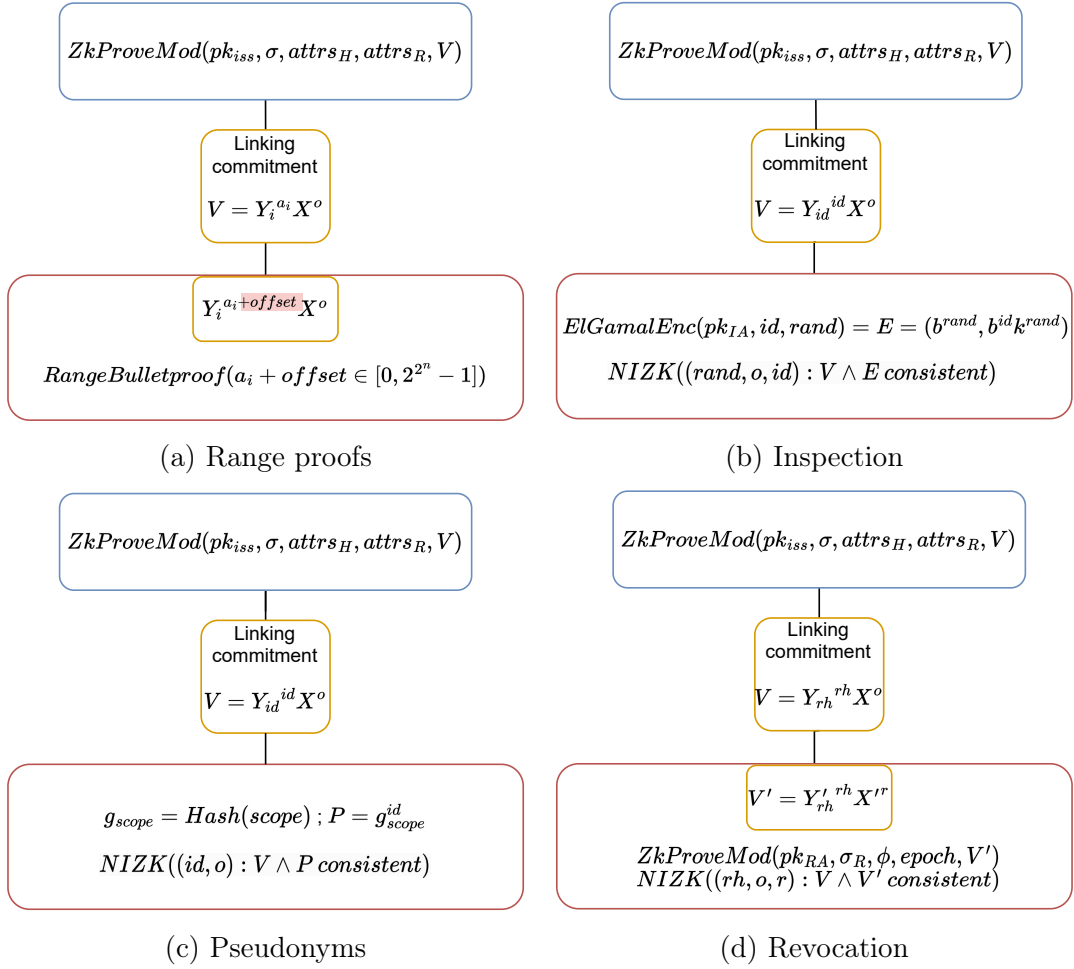


Figure 2.5: Schematic view of the applications implemented through the commit-and-prove extension

specific use cases, such as demonstrated in [23]. On another note, we also measured the impact on token size, which is summarised in Table 2.1. While the overhead from adding each proof is not negligible, the magnitudes will stay in the range of few kilobytes for the whole token in almost all use cases, which is notably low for practical purposes.

We compared the method with Idemix [30] (cf. Figure 2.7), the only implementation in the literature achieving a similarly modular approach and supporting all four predicates. Our implementation outperforms Idemix significantly, even when comparing proofs with advanced predicates against the basic policy. Additionally, the feasibility of the solution is shown through further scenarios. First, the execution time for each method is tested on a mobile phone, with most operations taking less than a second, the exception being range proofs that go from a few seconds to more than ten seconds at worst. Second, we show a typical proving scenario where the cryptographic operations are implemented in the C language, obtaining 5 to 10 times faster execution times. Thus, when necessary for specific cases, operations may be carried out in a more efficient way, though partially sacrificing the portability and flexibility of the original implementation.

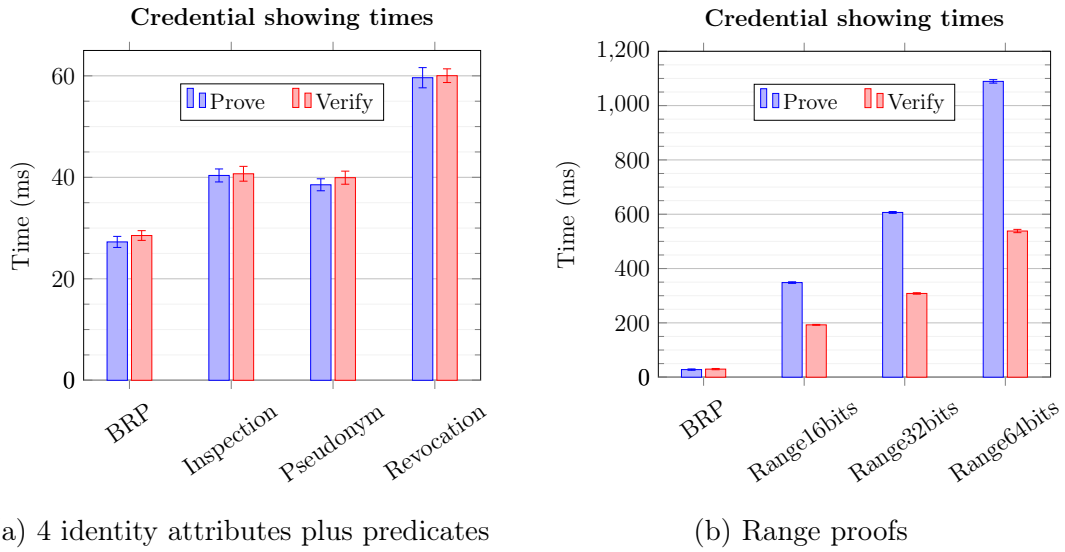


Figure 2.6: Execution times for proving and verifying when applying each of the predicates introduced along with the basic reveal policy as a reference point. For range proofs, 16-bit, 32-bit and 64-bit ranges are used

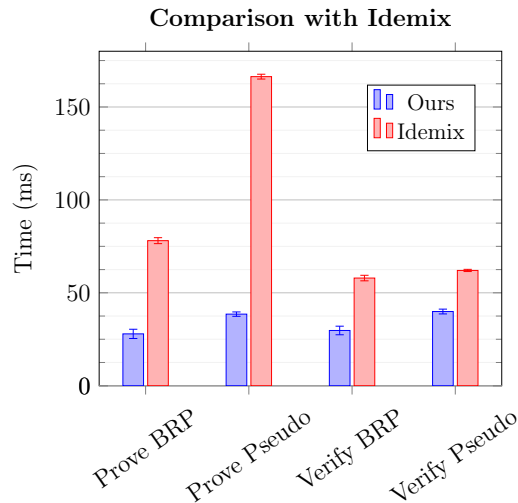


Figure 2.7: Comparison of Idemix with our implementation, both for a basic policy with statements for hiding/revealing attributes, and another with a pseudonym proof on top.

	Linking commitment	Pseudonym	Inspection	Revocation	Range 16-bit	Range 32-bit	Range 64-bit	4 hidden attributes	10 hidden attributes
Elements	$1\mathbb{G}_1+2\mathbb{Z}_p$	$1\mathbb{G}_1+3\mathbb{Z}_p$	$3\mathbb{G}_1+4\mathbb{Z}_p$	$2\mathbb{G}_1+10\mathbb{Z}_p$	$12\mathbb{G}_1+5\mathbb{Z}_p$	$14\mathbb{G}_1+5\mathbb{Z}_p$	$16\mathbb{G}_1+5\mathbb{Z}_p$	$2\mathbb{G}_1+9\mathbb{Z}_p$	$2\mathbb{G}_1+15\mathbb{Z}_p$
Bytes	193	241	483	674	1404	1598	1792	626	914

Table 2.1: Overhead in token size of the different predicates, along with the size increase of the proof because of a linking commitment, and token sizes for credential showings with 4 and 10 hidden attributes. Values correspond to using 48 and 97 bytes for  $\mathbb{Z}_p$  and  $\mathbb{G}_1$  element representation

All in all, we obtain an extensible, usable and practical implementation of dp-ABCs with advanced functionality. These characteristics (particularly when paired with the interoperability results discussed in this thesis) are very desirable in the current landscape of emerging SSI-based identity use cases. The implementation can then become a basis for introducing advanced privacy features into the existing heterogeneous scenarios, with a clear and simple way to adapt and improve the implementation whenever necessary. In this sense, Table 2.2 summarises how some contemporary use cases would benefit from the features achieved by the proposed solution.

Scenario	Features	Description
Identity wallets	All, extensibility	Current wallet efforts would benefit from unlinkability and zero-knowledge proofs. Generic heterogeneous scenarios and developing field make extensibility appealing.
IoT	Pseudonym, inspection	Scope-exclusive pseudonyms are commonly sought in IoT networks, while inspection enables trust management solutions.
Beyond 5G, 6G	Inspection, pseudonym, revocation	Privacy-preserving authentication to untrusted subnetworks. Pseudonyms improve efficiency (caching) and inspection/revocation enable security and accounting.
Biometric-based ABC	Range proof, extensibility	Range proofs enable linking a credential showing with a biometric matching operation. Extensibility allows adaptability of matching algorithm applied [23].
Green Pass	Range proof, revocation	The Green Pass requires both privacy features that can be met by advanced proofs like range proofs (e.g., on date of vaccination) and security guarantees for health safety (so revocation is important).
Academic degrees	Revocation	Enabling privacy-preserving usage of academic achievements while avoiding misuse through revocation [165].
Data spaces	Inspection, extensibility	Still-evolving frameworks with privacy-sensitive Attribute-Based Access and Usage Control. Inspection as an enabler of prevalent auditability processes.

Table 2.2: Overview of relevant use cases and the most appealing features of the proposed solution

### 2.2.3 To Pass or Not to Pass: Privacy-Preserving Physical Access Control

A rigorous formalisation of Biometric-Bound Attribute-Based Credentials (bb-ABC) was introduced in [23]. Moreover, it proposes two generic constructions which offer trade-offs between efficiency and trust assumptions, and provides micro-benchmarks showing the practical feasibility of a concrete instantiation using facial biometrics. Figure 2.8 shows an overview of the bb-ABC framework. It is comprised of the usual *issuers*, *users* and *verifiers* in p-ABC systems, complemented by a *reader device* in charge of scanning fresh user biometrics and generating from them input for users (U.GenEph and R.GenEphUser) and verifiers (V.InputGen and R.GenEphVerifier). This reader is an independent device deployed at verifier premises, avoiding dedicated hardware per-user to improve scalability. Thus, the device agglutinates the trust assumptions inherent to biometrics. Its reduced functionality and reusability for heterogeneous verifiers opens up straightforward certification and audit processes that make this trust affordable in practical scenarios. Further, the device will only require reduced communication and ephemeral cryptographic material, leading to a small attack surface.

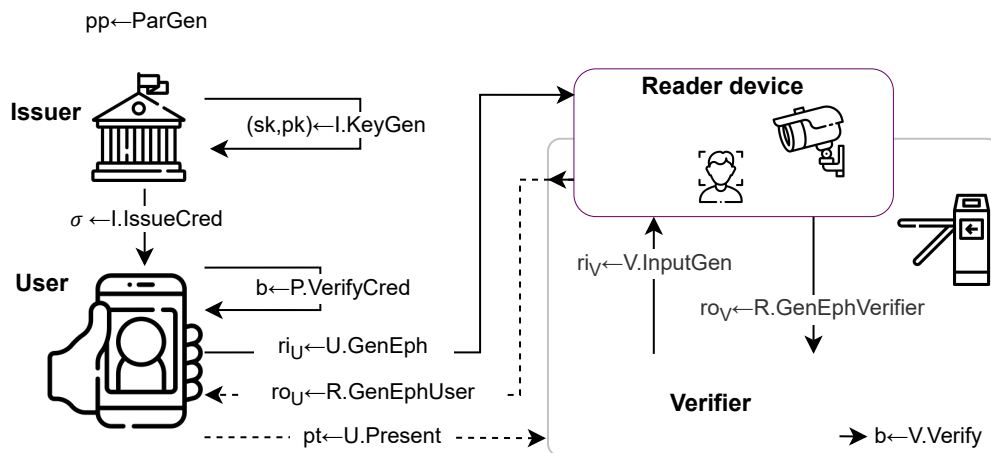


Figure 2.8: Biometric-Bound ABC components and communication flow

We propose a formal security model for the bb-ABC system, based on the usual notions for p-ABCs. In the following, we give an overview of the security properties, and refer to the full text of the paper for the formal definitions through adversarial experiments (cf. Section 4.3). **Correctness** follows the natural definition, with the particularity of being tied to the same false positive/negative rates as the biometric matching procedure implemented. **Unforgeability** requires that it is infeasible for an adversary to generate a valid presentation token as a forgery if it has not received a credential satisfying the predicates and the biometric matching, or has seen the exact same token. Before generating its output, the adversary is allowed to obtain any number of presentations of credentials of its choice, and request credentials on attributes at its discretion. This unforgeability notion directly covers non-transferability because presentations are bound to a biometric, which is specific to the credential owner. Lastly, **unlinkability** requires that no adversary

can link two user actions when the reader device is honestly executing its protocols. The adversary is given control over the issuer, the user's credential, and the verifier with the restriction that no trivial distinction may be possible. Notably, both definitions give the adversary full control over the biometrics used, which implies that a scheme proven secure in our model is also secure for any real-world distribution of biometrics.

Within this setting, we provide two generic constructions and prove them secure under the established model. In brief, the first construction, **BioABC-R**, defers the matching of biometric templates to the reader device. The user proves in zero-knowledge that a commitment of the template is consistent with the value embedded in the credential, and sends the plain template and opening values encrypted so that only the device can retrieve them for verifying biometric matching. This approach is largely agnostic to the specifics of the biometric method used.

In contrast, the **BioABC-ZK** construction relies on the user creating an actual proof in zero-knowledge of the biometric matching during credential presentation. In this way, the behaviour of the reader and verifier can be audited by users, and provably demonstrated to other parties. For example, this avoids the issue of a verifier arguing (fake) failed biometric matching to deny entry to a user, and helps detect incorrect or rogue behaviour by the reader. Thus, the necessary assumed trust in the overall system is reduced. However, as a trade-off the approach introduces a critical efficiency bottleneck for practical applications.

Indeed, the overhead over plain p-ABC showings in BioABC-R is simply an extra attribute (the hash of the biometric template), a symmetric authenticated encryption/decryption of the template, and a commitment opening check. Thus, it can be efficiently instantiated by state of art primitives such as PS-based p-ABCs, Pedersen commitments and authenticated AES-GCM-256 [166]. In contrast, the practicality of a BioABC-ZK instantiation will rest on the specifics of the biometric matching algorithm. For instance, state of art fingerprint matching is based on a procedural process for pairing minutiae [167], which is not translatable into efficient zero-knowledge proofs with existing techniques. Other biometric methods are currently hindered by complexity parameters instead. E.g., Iris recognition is based on computing the Hamming distance of biometric templates bigger than 10000 bits in practice. This leads to prohibitively high execution times with current hardware, although it can be an interesting approach in the near future.

Our instantiation, which is described in Construction 1 following the notation introduced in the publication [23], is based on matching facial biometrics. Particularly, we rely on Ouamane *et al.*'s [32] solution, that achieves 95% accuracy for "faces in the wild".<sup>9</sup> The method consists on extracting vector templates from biometric readings, and computing their cosine similarity: if the value is over a threshold, the templates are said to match. This process can be translated into an inner-product operation and a range proof, which can be efficiently computed in zero-knowledge. Particularly, we instantiate the range-proof as a simple bit-decomposition proof, as the prover will be more constrained than the verifier in our setting. The cryptographic primitives used in the instantiation are otherwise equivalent to the previous: p-ABC based on PS signatures, Pedersen Commitments and AES-GCM-256. In both instantiations, their security under the defined model is reduced to the security properties of the selected primitives. In fact, while long-term

<sup>9</sup><http://vis-www.cs.umass.edu/lfw/results.html>

<p><u>ParGen</u>(<math>1^\lambda</math>). Return <math>pp \leftarrow \\$ PS.\text{ParGen}(1^\lambda)</math></p> <p><b>Key Generation and Issuance.</b></p> <p><u>I.KeyGen</u>(<math>pp</math>). Return <math>(\text{sk}, \text{pk}) \leftarrow \\$ PS.\text{KeyGen}(pp)</math></p> <p><u>I.IssueCred</u>(<math>\text{sk}, a_{Bio}, \mathbf{a}</math>). <math>\mathbf{b} = (a_{Bio}, \mathbf{a})</math>. Return <math>\sigma \leftarrow \\$ PS.\text{Sign}(\text{sk}, \mathbf{b})</math></p> <p><u>U.VerifyCred</u>(<math>\text{pk}, \sigma, a_{Bio}, \mathbf{a}</math>). <math>\mathbf{b} = (a_{Bio}, \mathbf{a})</math>. Return 1 if <math>PS.\text{Verify}(\text{pk}, \mathbf{b}, \sigma) = 1</math>, else return 0</p> <p><b>Presentation.</b></p> <p><u>U.GenEph</u>(<math>pp</math>). Return <math>\text{sk}_{ae} \leftarrow \\$ AES.\text{KeyGen}()</math>.</p> <p><u>R.GenEphUser</u>(<math>Bio_f, ri_U</math>). Parse <math>ri_U = \{\text{sk}_{ae}\}</math>. Compute <math>(C_{Bio_f}, V_{Bio_f}) \leftarrow \\$ BitPC.\text{Commit}(Bio_f)</math>, i.e., <math>C_{Bio_f} = (C_1, \dots, C_N)</math> to the individual bits <math>f_i</math> of <math>Bio_f</math>, and <math>V_{Bio_f} = (r_1, \dots, r_N)</math> contains the individual openings. Return <math>ro_U = AES.\text{Encrypt}(\text{sk}_{ae}, \{C_{Bio_f}, V_{Bio_f}, Bio_f\})</math>.</p> <p><u>U.Present</u>(<math>\text{pk}, \sigma, a_{Bio}, \mathbf{a}, \phi, ro_U, ri_U, ctx</math>). Parse <math>\{C_{Bio_f}, V_{Bio_f}, Bio_f\} = AES.\text{Decrypt}(\text{sk}_{ae}, ro_U)</math>. If decryption fails, return <math>\perp</math>. <math>\mathbf{b} = (a_{Bio}, \mathbf{a})</math>.</p> <p>Given <math>a_{Bio} = (e_i)_{i \in [N]}</math>, <math>Bio_f = (f_i)_{i \in [N]}</math>. Choose random blinding values <math>w, z \leftarrow \\$ \mathbb{Z}_r</math>. Take <math>\sigma</math> as <math>(a', \sigma_1, \sigma_2)</math> and compute <math>(\sigma'_1, \sigma'_2) = (\sigma_1^w, (\sigma_2 \sigma_1^z)^w)</math>. Then, compute an Schnorr-style proof:  <math>pt \leftarrow \\$ NIZK[(\sigma, (e_i), \mathbf{a}, (f_i), (r_i), s, r) :</math></p> <p style="padding-left: 40px;">PS credential check <math>e(g_1^t X \prod_{i=1}^N Y_{e_i}^{e_i} \prod_{j=1}^n Y_{a_j}^{a_j} Y_{k+1}^{a'}, \sigma'_1) = e(g_1, \sigma'_2) e(X, \sigma'_1)^{-1} \wedge</math></p> <p style="padding-left: 40px;">Valid commitment <math>\bigwedge_{i=1}^N C_i = g^{f_i} h^{r_i} \wedge</math></p> <p style="padding-left: 40px;">Valid inner product <math>1 = \prod_{i=1}^N C_i^{e_i} g^{-s} h^{-r} \wedge</math></p> <p style="padding-left: 40px;">Biometric matches <math>s \in [2^{2l}\tau, 2^{2l}] \wedge</math></p> <p style="padding-left: 40px;">Predicate check <math>\phi(\mathbf{a}) = 1](\phi, ctx)</math></p> <p><b>Verification.</b></p> <p><u>V.InputGen</u>(<math>pt</math>). <math>ri_V = \varepsilon</math>.</p> <p><u>R.GenEphVerifier</u>(<math>Bio_f, ri_V</math>). Return <math>C_{Bio_f}</math> as computed in <u>R.GenEphUser</u></p> <p><u>V.Verify</u>(<math>\text{pk}, pt, \phi, ro_V, ctx</math>). Parse <math>C_{Bio_f} \leftarrow ro_V</math> and use it for verification of proof. If <math>pt</math> verifies correctly return 1. Else, return 0</p>
---

## Construction 1: Concrete instantiation of BioABC-ZKs

security is not achieved in terms of unforgeability, unlinkability is guaranteed even in post-quantum settings: Pedersen commitments are perfectly hiding, the NIZK is constructed through perfect honest-verifier zero-knowledge proofs and the Fiat-Shamir transform, and AES-GCM-256 preserves 128 bits of security even with attacks using Grover's algorithm.

We carried out micro-benchmarks for the cryptographic overhead introduced, showing the feasibility of the BioABC-ZK instantiation. Particularly, the overhead involves the

Entity	Process	Time (s)	Precomputable	Total (s)
<i>User</i>	PS cred	0.149	Yes	1.103
	Pedersen	0.463	No	
	Inner product	0.119	No	
	Bio match	0.372	Yes*	
<i>Verifier</i>	PS cred	0.060	No	0.415
	Pedersen	0.176	No	
	Inner product	0.044	No	
	Bio match	0.135	No	
<i>Reader</i>	Pedersen	1.677	Partially	1.677
<i>Total</i>	All processes without precomputation			3.195

Table 2.3: Computational overhead of BioABC-ZK instantiation

added proof goals during the zero-knowledge proof generation and verification, and the computation of the Pedersen commitment on the reader’s side. Each micro-benchmark was computed using a device representative to the setting: a mobile phone for the user, a general purpose laptop for the verifier and a Raspberry Pi 3 for the reader. The results are shown in Table 2.3, with a total of just over 3 seconds. Nonetheless, the implementation leaves room for additional optimisations, and particularly to precomputations that can reduce the online time. As reflected in the table, the user application can precompute the proof for the credential check and, depending on the setting, the biometric match proof. For its part, the reader can precompute the randomness in Pedersen commitments, which would halve the execution time. Just through the two precomputations, the execution time is lowered to around 2.1 seconds. The communication complexity is not prohibitive either, as the reader would need to send two messages of around 115KB and 58KB respectively, and the size of the proof sent by the user will be below 135KB.

## 2.2.4 A privacy-preserving attribute-based framework for IoT identity lifecycle management

In [19], a framework for privacy-preserving identity management in IoT scenarios is proposed. It comprehensively covers the challenges identified in the literature. Namely, establishing a root of trust for device identification (**Challenge-1**), using this root identity for bootstrapping in a domain (**Challenge-2**), enabling operational phase functionalities of identification (**Challenge-3**) and fine-granular authorisation while preserving privacy (**Challenge-4**), and achieving a coherent solution that links them (**Challenge-5**). The framework specifies steps at various phases during the device lifecycle. As shown in Figure 2.9, the initial bootstrapping will happen during or after device manufacturing, leading to the installation of *root identity* material with varying degrees of complexity according to device characteristics. Then, material related to the *security configuration* and characteristics of the device (i.e., *certificates*) are installed, potentially linked to the root identity in a cryptographic way. When the device is deployed in the operational domain, it performs a bootstrapping based on *authentication through the root identity*,

during which the security configurations and policies of the device will be adapted and applied to the security context. As the bootstrapping will ensue the creation of identification mechanisms for the device within the domain, the framework contemplates an optional procedure to register them through DIDs [44] following current SSI trends. Wrapping up the enrolment, the device has to execute an issuance process where it will receive domain *credentials* that attest its identity attributes. The process involves *identity proofs*, which may be tackled through the certificates obtained during the initial manufacturer bootstrapping, and can be adapted to the needs of the specific domain. After completing these steps, the device is ready to securely interact with other participants in the domain.

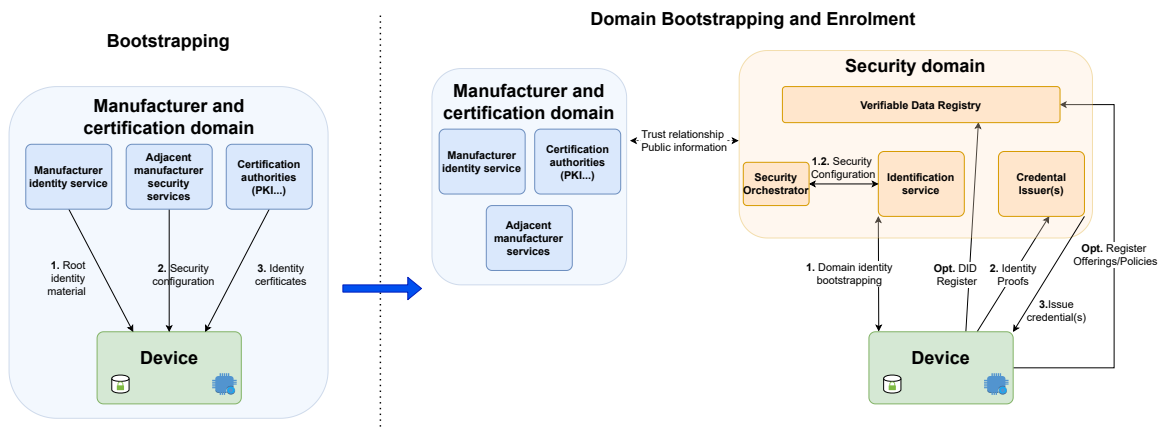


Figure 2.9: **Bootstrapping and enrolment** Components and steps for IoT device identity bootstrapping and enrolment phases

Particularly, the main goal of the framework is enabling fine-grained and privacy-preserving authorisation processes based on the zero-trust approach during the operational phase. As reflected in Figure 2.10, we consider two complementary approaches for coping with current solutions and heterogeneous scenarios. Both of them rely on attribute-based authentication through a *presentation token* generated by the device itself, which contains information requested by the required *policy*. While in Figure 2.10a, the service provider is in charge of checking the authorisation of the device, in Figure 2.10b this process is delegated to the infrastructure. At the cost of placing higher trust on centralised domain actors, the latter enables very constrained service providers to integrate with the advanced authorisation process. Moreover, it allows backwards compatibility to many scenarios that rely on infrastructure-based authorisation [156, 168]. Nevertheless, in most cases both approaches can coexist for a comprehensive coverage of the operational needs.

The framework was instantiated in the context of the EU H2020 ERATOSTHENES project, which aims to create and apply complementary identity, trust and lifecycle management techniques in IoT environments. In this instantiation, hardware-based PUFs [169] are used as a root of trust for device identification. After proper enrolment during the manufacturing phase, PUFs provide strong unclonability and unforgeability properties [170]. Additionally, the use of Trusted Execution Environments (TEE) is posed for protecting cryptographic material during its use. The security policy configurations for the device are established through an extended MUD file [16, 171]. Within the security domain, the



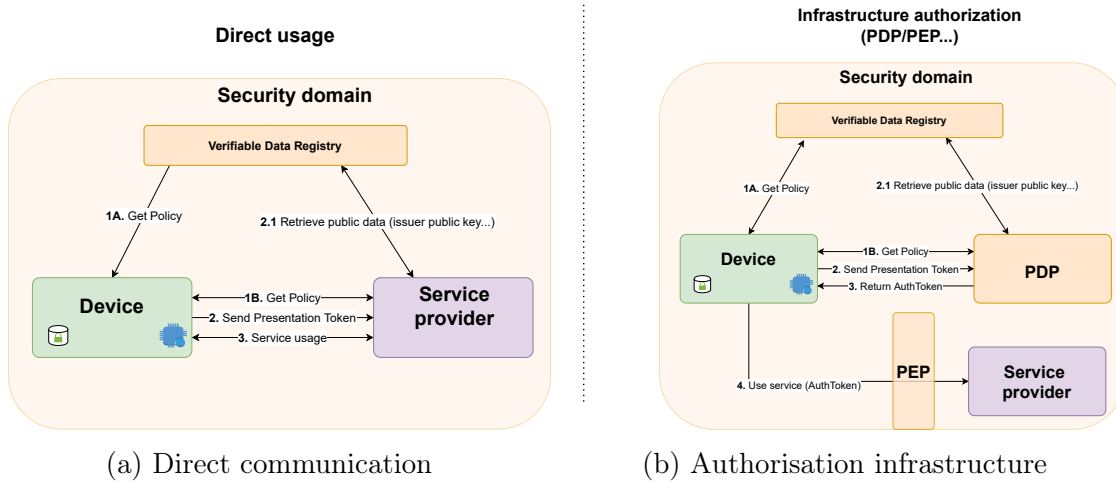


Figure 2.10: **Operational phase** Usage of the privacy-preserving approach based on attribute-based policies for authorisation in two alternative scenarios (a) one-to-one authorisation (b) authorisation through infrastructure

verifiable data registry is instantiated through DLTs, particularly Blockchain enhanced with smart contracts. For managing attribute-based identity, the VC specification [43] is used. Particularly, we apply a dp-ABC scheme [15, 18, 21] that enables unforgeability and unlinkability. The scheme is used for deriving zero-knowledge based presentations during authorisation processes, for which both the decentralised and centralised approaches are supported. Particularly, the infrastructure offers a Policy Enforcement Point (PEP) and Policy Decision Point (PDP) enhanced with capability-based authorisation tokens [172], which are secured using ECDSA-256 signatures, offering 128 bits of security [173].

Nonetheless, we note that flexibility is one of the main goals of the framework, as it is considered key to its practicality in heterogeneous IoT environments. Thus, variations over the main proposed instantiation are foreseen and planned. For instance, we envision that not every device will support PUFs, so alternative approaches for the root identity like pre-shared keys or traditional certificates may be considered. The collaboration with adjacent technologies is also considered. For example, the zero-trust approach of the framework helps accommodate such changes more smoothly, taking into account device characteristics to evaluate their trustworthiness.

In line with these ideas, we enhance the dp-ABC scheme to be aligned with the VC specification. Through this integration, we achieve transparent interchangeability between solutions with trade-offs on security, privacy, and efficiency, without changes to flows, implementations or models. The integration was developed within the context of the IoT application, driving its design. Namely, this lead to the following design goals: **C1-Trans**) The integration should be transparent. **C2-AdvZk**) It needs to support the use of advanced predicates over attributes. **C3-SecDom**) Identities will be used within specific security domains, and issuers will generate specific credential types tailored to them. **C4-IoT**) The application domain is an IoT environment. **C5-Unlink**) The technical solution must enable unlinkability.

The key point of integration is the definition of the “PsmsBlsSignature2022” signature

suite that handles VCs to generate and verify dp-ABC signatures, and the ‘PsmsBlsSignatureProof2022’ suite that serves to generate and verify a dp-ABC zero-knowledge proof over a VC. The role of these suites is to generate the input for the cryptographic operations, i.e., the identity attributes as used in the scheme. As a first step, it is necessary to perform normalisation of the VC so that syntactic differences such as serialisation order do not impede signing. In line with *C1-Trans*, we decided to follow current approaches and use the RDF Dataset Canonicalization Algorithm [174]. From the normalised output, the data to be signed by the cryptographic scheme is derived, that is, the list of values corresponding to the identity attributes. To allow predicates over attributes (*C2-AdvZk*), their parsing will depend on their type and characteristics like minimum and maximum value [20, 161]. Following the ideas in [20], we rely on *credential schemas* to define this information. For completeness, if no credential schema is present, all attributes are treated as plain strings that will only be considered for selective disclosure. Additionally, we opt for not signing metadata except for expiration date and the credential type and schema. This helps avoid privacy issues from excessive metadata (*C5-Unlink*), along with the impact on efficiency (*C4-IoT*) of extra attributes that should be kept hidden in most (if not all) authentication scenarios. Other metadata such as credential issuer may be included, but it will be implicitly verified and not signed as an attribute.

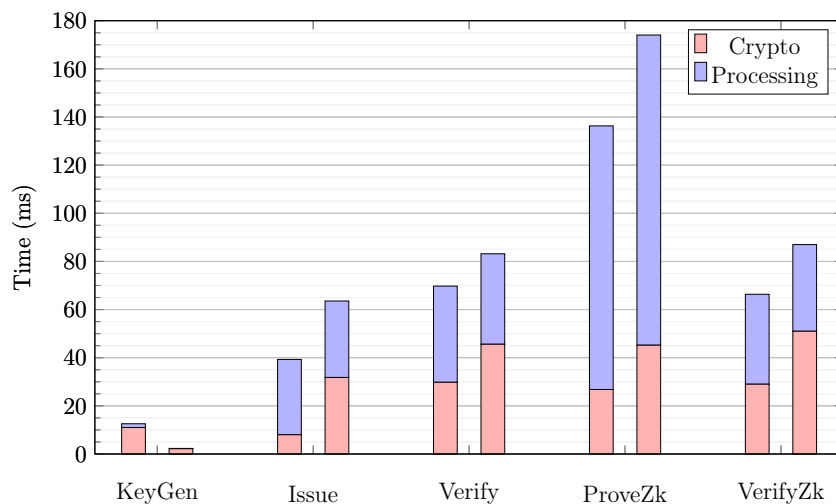


Figure 2.11: Execution time of the different phases in the Raspberry Pi 4 used as test device for *PS-MS* and *BBS+*, where *Crypto* refers to time spent specifically in the cryptographic operations of the method and *WalletProcessing* to all other operations in the wallet methods like VC parsing

Apart from these definitions, some other elements have to be determined to complete the integration. First, we designate the ‘Bls12381G1Key2022’ key format as a verification method. Apart from the mandatory fields, it contains the number of bases (and attributes supported) of the key. Note that, according to *C3-SecDom*, this will not result in an explosion of the number of keys needed.

On another note, we need to establish the format of the proofs themselves. Aligning with the *C1-Trans* condition, we decide to rely exclusively on the standard fields of the

specification. The definition of particular aspects of the proof, like the predicates applied, is left to the serialised proof value processed by the scheme. We present an alternative approach in [20], explicitly including the predicates as part of the derived credential. This would produce data more apparent for humans or machines processing the data, but it makes credentials, schemas, and parsing more complex.

Work	Challenge-1	Challenge-2	Challenge-3	Challenge-4	Challenge-5
This work	Yes	Yes	Yes	Yes	Yes
FDO [133]	Yes	Yes	Yes, no privacy	No	No
Choudhury <i>et al.</i> [148]	Yes	Yes	Yes	No	No
Canovas <i>et al.</i> [150, 151]	No	No	Yes	Yes	No
Alcaide <i>et al.</i> [152]	No	No	Yes	Yes	No
Niya <i>et al.</i> [134]	Yes	Yes	Yes, no privacy	No	No
Qureshi <i>et al.</i> [135]	Yes	Yes	Yes, no privacy	No	No
Kortesniemi <i>et al.</i> [136] and others [137–139]	No	No	Yes, no privacy	No	No
Diego <i>et al.</i> [140]	No	Yes	Yes, no privacy	Yes, no privacy	No
Akil <i>et al.</i> [149]	No	No	Yes	No	No

Table 2.4: Overview on state-of-art solutions on IoT identity management and which challenges they focus on tackling

We evaluated the performance of the implementation in a Raspberry 4 as a representative of high-end IoT devices, using BLS12-381 as the underlying curve. The results are shown in Figure 2.11, which also includes a comparison with BBS+. Our implementation slightly outperforms BBS+, though both approaches obtain practical values, with less than 150 ms for the costliest operation. Interestingly, the processing of the wallet operations and Verifiable Credentials takes similar or even more time than the cryptographic operations themselves.

In definitive, as summarised in Table 2.4, the proposed solution is the first to tackle the identified challenges for privacy preserving IoT identity management in a comprehensive manner. Nonetheless, other solutions, such as FDO [133], can be useful as complementary or alternative approaches, and future works may bring improvements to each of the identified phases. This extensibility is enabled by design through the generic approach and flexibility of the framework.

This chapter has given an overview on the outcomes achieved during the thesis development. The results have been developed within the framework of the identified challenges and gaps in the literature after exhaustive analysis [17–21], according to *Objective 1*. As part of the outcomes, we have proposed and evaluated various advances on privacy-preserving cryptographic systems, particularly in the field of distributed privacy-preserving Attribute-Based Credentials, fulfilling *Objective 2* [18, 21, 23]. Furthermore, in line with *Objectives 3 and 4*, the developed solutions have been integrated with emerging technologies such as DLTs or Verifiable Credentials, achieving comprehensive authentication and authorisation solutions in the scope of zero-trust architectures and Self-Sovereign

Identity [17, 19, 20, 22]. These developments have been extended to cover the needs of IoT environments, such as efficiency or adaptability to the heterogeneous characteristics of IoT contexts, achieving privacy-preserving identity management solutions encompassing devices lifecycle [19, 21, 23, 24] as outlined in *Objective 5*. Lastly, the practicality of these applications has been validated through various real-world use cases in the context of three H2020 European projects [17–19, 24–27], realizing *Objective 6*.

# Chapter 3

## Conclusions and future works

The widespread surge of digitisation of services, online presence of users and connectivity of IoT devices is shaping modern society. This trend has led to numerous benefits in efficiency, comfort and quality of life. Nonetheless, it also leads to a large attack surface for compromising security and privacy. In that sense, one of the main challenges is enabling secure authentication and authorisation through digital identity management while ensuring that individual's privacy rights are respected. In the current landscape, numerous use cases require such notions with heterogeneous constraints depending on the application scenario, as exemplified by Table 2.2. Existing solutions have failed to address gaps on security and privacy, such as the identity provider becoming a single-point of failure, and present issues that hamper their applicability in practical scenarios such as poor efficiency, extensibility or interoperability.

This thesis has been carried out with the goal of developing identity management solutions that are secure, privacy-preserving and practical, filling the gaps found in the literature. The main object of focus has been privacy-preserving Attribute-Based Credentials. We have prioritised the implementation and evaluation of distributed systems based on dp-ABCs that mitigate single points of failure. The implemented solution provides advanced capabilities such as range proofs or pseudonyms in a modular and extensible way through the commit-and-prove paradigm. Additionally, tackled the issue of non-transferability in p-ABCs through a biometric-based approach. We formalised the notion of Biometric-Bound Attribute-Based Credentials and proposed two practical instantiations, demonstrating their efficacy in real-world physical access control contexts.

Moreover, recognizing the importance of interoperability and standardisation for encouraging widespread adoption, we have integrated p-ABCs with the W3C Verifiable Credentials specification. This addresses one of the main shortcomings of existing p-ABCs, while also tackling the privacy concerns of the specification, which is seeing extensive use in current SSI proposals. Our aims for flexible, secure and user centric solutions have not been only focused on such technical outcomes, but in achieving comprehensive solutions to the problems at hand. Thus, we developed architectures based on zero-trust for authentication and authorisation with dp-ABCs as a key enabler of privacy and security, and Distributed Ledger Technologies as a backbone for decentralised trust.

Furthermore, we have emphasised the importance of extending efforts for achieving

security and privacy in identity management to encompass IoT environments, given their pervasive nature and unique challenges. In this scope, we proposed and instantiated a framework that spans the whole lifecycle of IoT devices. The framework provides a way to realise trustworthy and privacy-preserving authorisation during the operation within IoT security contexts. This work underscores the significance of flexibility to accommodate the heterogeneity of these environments.

Lastly, we have been committed to the practical validation of the solutions proposed within this thesis in relevant use cases. This has resulted in their exploitation as part of three Horizon 2020 European projects (OLYMPUS, CyberSec4Europe and ERATOS-THENES) demonstrating their feasibility. Summarizing:

- Distributed privacy-preserving Attribute-Based Credentials (dp-ABC) have been applied to address identity management in an efficient and usable way while mitigating the single point of failure issue of the issuer.
- The use of dp-ABC enables privacy, security and functionality notions inherent to ecosystems based on zero-trust authorisation.
- The leveraged presentation process of p-ABCs based on linking commitments and commit-and-prove techniques enables modular extensibility retaining formal guarantees of security and privacy.
- The p-ABC extension was demonstrated as a mechanism to efficiently achieve key advanced p-ABC features for contemporary identity use cases, namely inspection, pseudonyms, revocation and range proofs.
- The concept of Biometric-Bound Attribute-Based Credentials (bb-ABC), which enables non-transferability through biometrics without dedicated devices per user, has been formalised through the definition of variants of the traditional p-ABC security properties: correctness, soundness and unlinkability.
- The practical relevance of bb-ABC was demonstrated through two complementary constructions and their instantiation with concrete primitives, whose efficiency was showcased through micro-benchmarks.
- The applicability of dp-ABC in relevant use cases, particularly IoT environments, was improved through their effective integration into the W3C Verifiable Credential specification.
- A framework for flexible, comprehensive and privacy-preserving identity management of IoT devices throughout their lifecycle is crucial, and has been achieved with dp-ABCs as a key enabler.
- The developed solutions were successfully applied to solve challenges in multiple relevant use cases, showcasing their relevance in the current identity landscape.

## 3.1 Future work

From these results, a number of interesting avenues for continued research can be identified.

First, the developed dp-ABC system enables future endeavors by itself. Its extensibility opens paths of improvement like adding blind issuance of attributes or additional proofs such as set-membership. Any theoretical work in related building blocks, such as commit-and-prove techniques for, e.g., additional revocation mechanisms based on accumulators would also translate into direct benefits. Regarding the non-transferability approach, a full deployment of the system and the exploration of additional biometric features and matching algorithms that might become p-ABC-friendly could be pursued. Additionally, the W3C's Verifiable Credentials specification itself and its integration with p-ABCs, could be further improved to enable the flexibility and parametrisation of signature suites, further leaning into the strong points of this result. A particularly impactful outcome in this line of work could be derived from linking with current efforts to achieve European Digital Identity through identity wallets. The initial deployments will not include the advanced privacy properties offered by p-ABCs. Our solution may be integrated as an extension of the initial wallet features, empowering citizens' control over their privacy rights and services offered by identity providers. An additional challenge is the attestation of the wallet itself in a privacy-respecting way. This innovation is being pursued in further European and national projects such as LICORICE and TruIdentity.

On another note, the developed generic IoT identity management framework can be instantiated through other state-of-art tools, serving to explore different scenarios and their appropriateness for tackling specific challenges or developing adaptations to the processes. In this sense, the identity management framework could be extended to cover the needs of ultra-low-end devices through symmetric primitives in an integrated solution. In this direction, slightly veering from identity management itself, the achievement of fully functional security profiles and targets would be a key step for achieving completely flexible identity management approaches and other cybersecurity aspects.

Lastly, as the general trend in the realm of cryptography, one may explore the achievement of equivalent or improved security and privacy solutions in a post-quantum scenario. Particularly, in line with the thesis' focus on flexibility and configurability of solutions, we envision a line of research and innovation that aims to ease the transition between traditional cryptography and post-quantum solutions through their complementary application until the former can be phased out.

# Chapter 4

## Publications composing the doctoral thesis

### 4.1 Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures

**Abstract** Despite the latest efforts to foster the adoption of privacy-enhancing Attribute-Based Credential (p-ABC) systems in electronic services, those systems are not yet broadly adopted mainly due to performance efficiency issues, lack of interoperability with standards and the centralized architectural scheme that relies on a unique Identity Provider (IdP) for credential issuance. To cope with these limitations, this paper describes the first implementation of the Pointcheval-Sanders Multi-Signatures (PS-MS) crypto scheme proposed by Camenisch et al. and its integration in a distributed and privacy-preserving identity management system proposed in OLYMPUS H2020 European research project. Our efficient implementation provides remarkable privacy-preservation features for identity management in online transactions leveraging p-ABC systems, including unforgeability, minimal disclosure of personal data through zero-knowledge proofs, unlinkability in online transactions and fully distributed credential issuance across different IdPs, thereby removing the IdP as a unique point of failure. The performance of the implementation has been exhaustively analyzed and evaluated with different curves, signers and number of attributes, and compared against Identity Mixer, the best known p-ABC system, outperforming significantly the credential issuance and zero-knowledge proving and verification processes (2-4 times less execution time).



<b>Title</b>	Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures
<b>Authors</b>	Jesús García-Rodríguez, Rafael Torres Moreno, Jorge Bernal Bernabé, Antonio Skarmeta
<b>Journal</b>	Journal of Information Security and Applications
<b>Impact factor (2021)</b>	4.960
<b>JCR Rank (2021)</b>	Computer science, information systems: 44/164
<b>Publisher</b>	Elsevier
<b>Date</b>	November 2021
<b>ISSN</b>	2214-2126
<b>EISSN</b>	2214-2134
<b>DOI</b>	<a href="https://doi.org/10.1016/j.jisa.2021.102971">https://doi.org/10.1016/j.jisa.2021.102971</a>
<b>State</b>	Published
<b>Contribution</b>	Methodology, Software, Writing – original draft, Investigation, Formal analysis

## 4.2 A privacy-preserving attribute-based framework for IoT identity lifecycle management

**Abstract** The Internet of Things (IoT) has brought a new era of interconnected devices and seamless data exchange. As the IoT ecosystem continues to expand, there is an increasing need for effective identity management mechanisms, specifically for authorization processes and access control. The pervasiveness of such devices demands that desirable solutions tackle not only security properties but also privacy aspects like granular control over which identity data is shared in authentication/authorization processes, covering aspects like bootstrapping, enrolment, and service provision. In this context, it is natural to turn to privacy-enhancing technologies, like (privacy-preserving) Attribute-Based Credentials (p-ABC), for achieving both high security and privacy guarantees. Nonetheless, these technical tools need to be accompanied by a comprehensive approach that deals with the particularities of IoT scenarios and covers the full lifetime of the device. In this work, we propose the use of a p-ABC scheme with support for distributed issuance (dp-ABC) as a keystone for privacy-preserving attribute-based authentication and authorization in IoT scenarios. We integrate said cryptographic scheme with W3C’s Verifiable Credentials standard, evaluating its impact to gauge its feasibility. The integration facilitates adoption and, particularly, allows the solution to transparently coexist with simpler techniques in heterogeneous scenarios that demand them. Moreover, we define and analyse a generic and comprehensive framework for identity management that identifies challenges throughout the device’s lifetime to achieve IoT privacy-preserving identity management following self-sovereign principles. We show how the various aspects identified in the framework are tackled in a concrete instantiation as part of the H2020 project ERATOSTHENES.

<b>Title</b>	A privacy-preserving attribute-based framework for IoT identity lifecycle management
<b>Authors</b>	Jesús García-Rodríguez, Antonio Skarmeta
<b>Journal</b>	Computer Networks
<b>Impact factor (2022)<sup>3</sup></b>	5.6
<b>JCR Rank (2022)<sup>3</sup></b>	Computer science, hardware & architecture: 8/54 Computer science, information systems: 41/158
<b>Publisher</b>	Elsevier
<b>Date</b>	November 2023
<b>ISSN</b>	1389-1286
<b>EISSN</b>	1872-7069
<b>DOI</b>	<a href="https://doi.org/10.1016/j.comnet.2023.110039">https://doi.org/10.1016/j.comnet.2023.110039</a>
<b>State</b>	Published
<b>Contribution</b>	Methodology, Software, Formal analysis, Writing – original draft, Writing – review & editing, Visualization, Investigation

### 4.3 To pass or not to pass: Privacy-preserving physical access control

**Abstract** Anonymous or attribute-based credential (ABC) systems are a versatile and important cryptographic tool to achieve strong access control guarantees while simultaneously respecting the privacy of individuals. A major problem in the practical adoption of ABCs is their transferability, i.e., such credentials can easily be duplicated, shared or lent. One way to counter this problem is to tie ABCs to biometric features of the credential holder and to require biometric verification on every use. While this is certainly not a viable solution for all ABC use-cases, there are relevant and timely use-cases, such as vaccination credentials as widely deployed during the COVID-19 pandemic. In such settings, ABCs that are tied to biometrics, which we call Biometric-Bound Attribute-Based Credentials (bb-ABC), allow to implement scalable and privacy-friendly systems to control physical access to (critical) infrastructure and facilities.

While there are some previous works on bb-ABC in the literature, the state of affairs is not satisfactory. Firstly, in existing work the problem is treated in a very abstract way when it comes to the actual type of biometrics. Thus, it does not provide concrete solutions which allow for assessing their practicality when deployed in a real-world setting. Secondly, there is no formal model which rigorously captures bb-ABC systems and their security requirements, making it hard to assess their security guarantees. With this work we overcome these limitations and provide a rigorous formalization of bb-ABC systems. Moreover, we introduce two generic constructions which offer different trade-offs between efficiency and trust assumptions, and provide benchmarks from a concrete instantiation of such a system using facial biometrics. The latter represents a contact-less biometric feature that provides acceptable accuracy and seems particularly suitable to the above use-case.

<b>Title</b>	To pass or not to pass: Privacy-preserving physical access control
<b>Authors</b>	Jesús García-Rodríguez, Stephan Krenn, Daniel Slamanig
<b>Journal</b>	Computers & Security
<b>Impact factor (2022)<sup>3</sup></b>	5.6
<b>JCR Rank (2022)<sup>3</sup></b>	Computer science, information systems: 41/158
<b>Publisher</b>	Elsevier
<b>Date</b>	January 2024
<b>ISSN</b>	0167-4048
<b>EISSN</b>	1872-6208
<b>DOI</b>	<a href="https://doi.org/10.1016/j.cose.2023.103566">https://doi.org/10.1016/j.cose.2023.103566</a>
<b>State</b>	Published
<b>Contribution</b>	Conceptualization, Formal analysis, Funding acquisition, Software, Writing – original draft, Writing – review & editing

<sup>3</sup>At the time of writing of this thesis, the impact factor and JCR rank data for the year 2023/2024 were not available.

## 4.4 Beyond Selective Disclosure: Extending Distributed p-ABC Implementations by Commit-and-Prove Techniques

**Abstract** The increasing user awareness and regulatory framework (e.g., GDPR, eIDAS2) have contributed to considering data minimization and privacy-by-design as central guiding principles for new systems. Among others, this has led to a paradigm shift towards Self-Sovereign Identity solutions to put the user in full control over their data. Despite the promising landscape, privacy-preserving Attribute-Based Credentials (p-ABC) have not been widely adopted, mainly due to the lack of secure, flexible and efficient implementations that cover the basic and advanced needs in p-ABC systems. In this work, we tackle this gap by developing an improved zero-knowledge showing protocol of a distributed p-ABC scheme based on Pointcheval-Sanders Multi-Signatures to allow for modular extensions through commit-and-prove techniques. We use it to implement a flexible p-ABC system with decentralized issuance that, apart from the basic notions of p-ABCs, covers range proofs, pseudonyms, inspection and revocation. Lastly, we thoroughly evaluate the performance of the system under different testbed conditions, showing a significant efficiency improvement over previous implementations.

<b>Title</b>	Beyond Selective Disclosure: Extending Distributed p-ABC Implementations by Commit-and-Prove Techniques
<b>Authors</b>	Jesús García-Rodríguez, Stephan Krenn, Jorge Bernal Bernabé, Antonio Skarmeta
<b>Journal</b>	Computer Networks
<b>Impact factor (2022)<sup>3</sup></b>	5.6
<b>JCR Rank (2022)<sup>3</sup></b>	Computer science, hardware & architecture: 8/54 Computer science, information systems: 41/158
<b>Publisher</b>	Elsevier
<b>Date</b>	June 2024
<b>ISSN</b>	1389-1286
<b>EISSN</b>	1872-7069
<b>DOI</b>	<a href="https://doi.org/10.1016/j.comnet.2024.110498">https://doi.org/10.1016/j.comnet.2024.110498</a>
<b>State</b>	Published
<b>Contribution</b>	Methodology, Software, Formal analysis, Writing – original draft, Writing – review & editing, Visualization, Investigation



# Bibliography

- [1] C. Forrest, 2012 dropbox hack worse than realized, 68m passwords leaked (2016).  
URL <https://www.techrepublic.com/article/2012-dropbox-hack-worse-than-realized-68m-passwords-leaked/>
- [2] D. Kocieniewski, Adobe announces security breach (2013).  
URL <https://www.nytimes.com/2013/10/04/technology/adobe-announces-security-breach.html>
- [3] A. Scroxton, Leak of 26 billion records may prove to be ‘mother of all breaches’ (2024).  
URL <https://www.computerweekly.com/news/366567105/Leak-of-26-billion-records-may-prove-to-be-mother-of-all-breaches>
- [4] A. Forrest, Facebook data scandal: Social network fined \$5bn over ‘inappropriate’ sharing of users’ personal information (2016).  
URL <https://www.independent.co.uk/news/world/americas/facebook-data-privacy-scandal-settlement-cambridge-analytica-court-a9003106.html>
- [5] K. Ashton, et al., That ‘internet of things’ thing, RFID journal 22 (7) (2009) 97–114.
- [6] L. Vailshery, Number of internet of things (iot) connected devices worldwide from 2019 to 2023, with forecasts from 2022 to 2030 (2024).  
URL <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/>
- [7] A. Tobin, D. Reed, The inevitable rise of self-sovereign identity, The Sovrin Foundation (2016).  
URL <https://sovrin.org/wp-content/uploads/2017/06/The-Inevitable-Rise-of-Self-Sovereign-Identity.pdf>
- [8] European Digital Identity - Provisional Agreement | Highlights | Home | ITRE | Committees | European Parliament — europarl.europa.eu, <https://www.europarl.europa.eu/committees/en/european-digital-identity-provisional-ag/product-details/20231116CAN72103>, [Accessed 18-12-2023] (2023).
- [9] Sovrin Foundation, Self-sovereign identity and iot (2020).  
URL <https://identity.foundation/bbs-signature/draft-irtf-cfrg-bbs-signatures>

- 
- [10] D. Chaum, Untraceable electronic mail, return addresses, and digital pseudonyms, *Commun. ACM* 24 (2) (1981) 84–88. doi:10.1145/358549.358563.  
URL <https://doi.org/10.1145/358549.358563>
- [11] D. Chaum, Security without identification: Transaction systems to make big brother obsolete, *Commun. ACM* 28 (10) (1985) 1030–1044. doi:10.1145/4372.4373.
- [12] Olympus: Oblivious identity management for private and user-friendly services, H2020-DS-2016-2017 (2018).  
URL <https://doi.org/10.3030/786725>
- [13] Cybersec4europe: Cyber security network of competence centres for europe, H2020-DS-2016-2017 (2019).  
URL <https://doi.org/10.3030/830929>
- [14] Eratosthenes: Secure management of iot devices lifecycle through identities, trust and distributed ledgers, H2020-SU-DS-2018-2019-2020 (2021).  
URL <https://doi.org/10.3030/101020416>
- [15] J. Camenisch, M. Drijvers, A. Lehmann, G. Neven, P. Towa, Short threshold dynamic group signatures, in: *International Conference on Security and Cryptography for Networks*, Springer, 2020, pp. 401–423.
- [16] E. Lear, R. Droms, D. Romascanu, Manufacturer usage description specification, Tech. rep., RFC Editor (2019).
- [17] R. Torres Moreno, J. Bernal Bernabe, J. García Rodríguez, T. Kasper Frederiksen, M. Stausholm, N. Martínez, E. Sakkopoulos, N. Ponte, A. Skarmeta, The OLYMPUS Architecture—Oblivious Identity Management for Private User-Friendly Services, *Sensors* 20 (3) (2020) 945. doi:10.3390/s20030945.  
URL <https://www.mdpi.com/1424-8220/20/3/945>
- [18] J. García-Rodríguez, R. T. Moreno, J. B. Bernabé, A. F. Skarmeta, Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures, *Journal of Information Security and Applications* 62 (2021) 102971. doi:10.1016/j.jisa.2021.102971.  
URL <https://doi.org/10.1016/j.jisa.2021.102971>
- [19] J. García-Rodríguez, A. Skarmeta, A privacy-preserving attribute-based framework for IoT identity lifecycle management, *Computer Networks* 236 (2023) 110039. doi:10.1016/j.comnet.2023.110039.  
URL <https://linkinghub.elsevier.com/retrieve/pii/S138912862300484X>
- [20] J. García-Rodríguez, R. T. Moreno, J. B. Bernabé, A. F. Skarmeta, Towards a standardized model for privacy-preserving verifiable credentials, in: D. Reinhardt, T. Müller (Eds.), *ARES 2021: The 16th International Conference on Availability, Reliability and Security*, Vienna, Austria, August 17-20, 2021, ACM, 2021, pp. 126:1–126:6. doi:10.1145/3465481.3469204.  
URL <https://doi.org/10.1145/3465481.3469204>

- 
- [21] J. García-Rodríguez, S. Krenn, J. Bernal Bernabe, A. Skarmeta, Beyond selective disclosure: Extending distributed p-ABC implementations by commit-and-prove techniques, *Computer Networks* 248 (2024) 110498. doi:10.1016/j.comnet.2024.110498. URL <https://linkinghub.elsevier.com/retrieve/pii/S138912862400330X>
- [22] S. Daoudagh, E. Marchetti, V. Savarino, J. B. Bernabe, J. García-Rodríguez, R. T. Moreno, J. A. Martinez, A. F. Skarmeta, Data Protection by Design in the Context of Smart Cities: A Consent and Access Control Proposal, *Sensors* 21 (21) (2021) 7154. doi:10.3390/s21217154. URL <https://www.mdpi.com/1424-8220/21/21/7154>
- [23] J. García-Rodríguez, S. Krenn, D. Slamanig, To pass or not to pass: Privacy-preserving physical access control, *Computers & Security* 136 (2024) 103566. doi:10.1016/j.cose.2023.103566. URL <https://linkinghub.elsevier.com/retrieve/pii/S0167404823004765>
- [24] K. Loupos, H. Niavis, F. Michalopoulos, G. Misiakoulis, A. Skarmeta, J. Garcia, A. Palomares, H. Song, R. Dautov, F. Giampaolo, R. Mancilla, F. Costantino, D. Van Landuyt, S. Michiels, S. More, C. Xenakis, M. Bampatsikos, I. Politis, K. Krilakis, S. Vavilis, An inclusive Lifecycle Approach for IoT Devices Trust and Identity Management, in: *Proceedings of the 18th International Conference on Availability, Reliability and Security*, ACM, Benevento Italy, 2023, pp. 1–6. doi:10.1145/3600160.3605083. URL <https://dl.acm.org/doi/10.1145/3600160.3605083>
- [25] R. T. Moreno, J. G. Rodriguez, C. T. Lopez, J. B. Bernabe, A. Skarmeta, OLYMPUS: A distributed privacy-preserving identity management system, in: *2020 Global Internet of Things Summit (GIoTS)*, IEEE, Dublin, Ireland, 2020, pp. 1–6. doi:10.1109/GIoTTS49054.2020.9119663. URL <https://ieeexplore.ieee.org/document/9119663/>
- [26] J. B. Bernabe, J. García-Rodríguez, S. Krenn, V. Liagkou, A. Skarmeta, R. Torres, Privacy-Preserving Identity Management and Applications to Academic Degree Verification, in: M. Friedewald, S. Krenn, I. Schiering, S. Schiffner (Eds.), *Privacy and Identity Management. Between Data Protection and Security*, Vol. 644, Springer International Publishing, Cham, 2022, pp. 33–46, series Title: IFIP Advances in Information and Communication Technology. doi:10.1007/978-3-030-99100-5\_4. URL [https://link.springer.com/10.1007/978-3-030-99100-5\\_4](https://link.springer.com/10.1007/978-3-030-99100-5_4)
- [27] J. García-Rodríguez, D. Goodman, S. Krenn, V. Liagkou, R. T. Moreno, From Research to Privacy-Preserving Industry Applications: Workshop Summary, in: F. Bieker, J. Meyer, S. Pape, I. Schiering, A. Weich (Eds.), *Privacy and Identity Management*, Vol. 671, Springer Nature Switzerland, Cham, 2023, pp. 21–33, series Title: IFIP Advances in Information and Communication Technology. doi:10.1007/978-3-031-31971-6\_3. URL [https://link.springer.com/10.1007/978-3-031-31971-6\\_3](https://link.springer.com/10.1007/978-3-031-31971-6_3)



- 
- [28] J. Camenisch, A. Lysyanskaya, A signature scheme with efficient protocols, in: S. Cimato, C. Galdi, G. Persiano (Eds.), Security in Communication Networks, Third International Conference, SCN 2002, Amalfi, Italy, September 11-13, 2002. Revised Papers, Vol. 2576 of Lecture Notes in Computer Science, Springer, 2002, pp. 268–289.
- [29] A. Sabouri, K. Rannenberg, Abc4trust: protecting privacy in identity management by bringing privacy-abcs into real-life, in: IFIP International Summer School on Privacy and Identity Management, Springer, 2014, pp. 3–16.
- [30] J. B. Bernabe, M. David, R. T. Moreno, J. P. Cordero, S. Bahloul, A. Skarmeta, Aries: Evaluation of a reliable and privacy-preserving european identity management framework, Future Generation Computer Systems 102 (2020) 409–425.
- [31] T. P. Pedersen, Non-interactive and information-theoretic secure verifiable secret sharing, in: CRYPTO 1991, Springer, 1991, pp. 129–140.
- [32] A. Ouamane, M. Bengherabi, A. Hadid, M. Cheriet, Side-information based exponential discriminant analysis for face verification in the wild, in: 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition, FG 2015, Ljubljana, Slovenia, May 4-8, 2015, IEEE Computer Society, 2015, pp. 1–6. doi:10.1109/FG.2015.7284837.  
URL <https://doi.org/10.1109/FG.2015.7284837>
- [33] ISO Central Secretary, Information technology – security techniques – guidance for the production of protection profiles and security targets, Standard ISO/IEC TR 15446:2017, International Organization for Standardization, Geneva, CH (2017).  
URL <https://www.iso.org/standard/68904.html>
- [34] OECD, Emerging privacy-enhancing technologies (2023). doi:10.1787/bf121be4-en.  
URL <https://www.oecd-ilibrary.org/content/paper/bf121be4-en>
- [35] C. Adams, S. Farrell, T. Kause, T. Mononen, Internet x. 509 public key infrastructure certificate management protocol (cmp), Tech. rep., RFC 4210 (Proposed Standard) (2005).
- [36] R. Perlman, An overview of PKI trust models, IEEE Network 13 (6) (1999) 38–43. doi:10.1109/65.806987.  
URL <http://ieeexplore.ieee.org/document/806987/>
- [37] A. Nash, H. Studiawan, G. Grispos, K.-K. R. Choo, Security analysis of google authenticator, microsoft authenticator, and authy, in: International Conference on Digital Forensics and Cyber Crime, Springer, 2023, pp. 197–206.
- [38] D. W. Chadwick, Federated identity management, in: Foundations of security analysis and design V, Springer, 2009, pp. 96–120.

- [39] B. Campbell, C. Mortimore, M. Jones, Security assertion markup language (saml) 2.0 profile for oauth 2.0 client authentication and authorization grants, RFC 7522, RFC Editor (May 2015).
- [40] D. Hardt, The oauth 2.0 authorization framework, Tech. rep., IETF (2012).
- [41] D. Recordon, D. Reed, Openid 2.0: a platform for user-centric identity management, in: Proceedings of the second ACM workshop on Digital identity management, ACM, 2006, pp. 11–16.
- [42] G. Eggers et al., Gaia-x technical architecture (2020).  
URL <https://www.data-infrastructure.eu/GAIAx/Redaktion/EN/Publications/gaia-x-technical-architecture.pdf>
- [43] M. Sporny, D. Longley, D. Chadwick, O. Steele, Verifiable credentials data model v2.0, W3C, W3C Recommendation (2024).  
URL <https://www.w3.org/TR/vc-data-model-2.0/>
- [44] D. Reed, M. Sporny, D. Longley, C. Allen, R. Grant, M. Sabadello, J. Holt, Decentralized identifiers (dids) v1. 0, Draft Community Group Report (2020).
- [45] H. Halpin, Vision: A Critique of Immunity Passports and W3C Decentralized Identifiers, in: T. Van Der Merwe, C. Mitchell, M. Mehrnezhad (Eds.), Security Standardisation Research, Vol. 12529, Springer International Publishing, Cham, 2020, pp. 148–168, series Title: Lecture Notes in Computer Science. doi:10.1007/978-3-030-64357-7\_7.  
URL [https://link.springer.com/10.1007/978-3-030-64357-7\\_7](https://link.springer.com/10.1007/978-3-030-64357-7_7)
- [46] C. Brunner, U. Gellersdörfer, F. Knirsch, D. Engel, F. Matthes, DID and VC:Untangling Decentralized Identifiers and Verifiable Credentials for the Web of Trust, in: 2020 the 3rd International Conference on Blockchain Technology and Applications, ACM, Xi’an China, 2020, pp. 61–66. doi:10.1145/3446983.3446992.  
URL <https://dl.acm.org/doi/10.1145/3446983.3446992>
- [47] D. Fett, K. Yasuda, B. Campbell, Selective disclosure for jwts (sd-jwt), Tech. rep., IETF Internet-Draft. (2024).
- [48] T. Lodderstedt, K. Yasuda, T. Looker, Openid for verifiable credential issuance, [https://openid.net/specs/openid-4-verifiable-credential-issuance-1\\_0.html](https://openid.net/specs/openid-4-verifiable-credential-issuance-1_0.html), accessed: 2024-04-03 (2024).
- [49] T. Lodderstedt, K. Yasuda, T. Looker, Openid for verifiable presentations, [https://openid.net/specs/openid-4-verifiable-presentations-1\\_0.html](https://openid.net/specs/openid-4-verifiable-presentations-1_0.html), accessed: 2024-04-03 (2023).
- [50] K. Yasuda, M. Jones, T. Lodderstedt, Self-issued openid provider v2, [https://openid.net/specs/openid-connect-self-issued-v2-1\\_0.html](https://openid.net/specs/openid-connect-self-issued-v2-1_0.html), accessed: 2024-04-03 (2023).

- 
- [51] Dsba releases ‘technical convergence discussion document’, <https://data-spaces-business-alliance.eu/dsba-releases-technical-convergence-discussion-document/>, accessed: 2024-04-04 (2023).
- [52] V. Dhillon, D. Metcalf, M. Hooper, The hyperledger project, in: Blockchain enabled applications, Springer, 2017, pp. 139–149.
- [53] G. Wood, et al., Ethereum: A secure decentralised generalised transaction ledger, Ethereum project yellow paper 151 (2014) (2014) 1–32.
- [54] D. Khovratovich, J. Law, Sovrin: digital identities in the blockchain era, Github Commit by jasonalaw October 17 (2017).
- [55] A decentralized, open source solution for digital identity and access management (2014).  
URL <https://github.com/jolocom>
- [56] W. F. Silvano, R. Marcelino, Iota tangle: A cryptocurrency to communicate internet-of-things data, Future generation computer systems 112 (2020) 307–319.
- [57] X. Chen, R. Nakada, K. Nguyen, H. Sekiya, A Comparison of Distributed Ledger Technologies in IoT: IOTA versus Ethereum, in: 2021 20th International Symposium on Communications and Information Technologies (ISCIT), IEEE, Tottori, Japan, 2021, pp. 182–187. doi:10.1109/ISCIT52804.2021.9590601.  
URL <https://ieeexplore.ieee.org/document/9590601/>
- [58] J. Bernal Bernabe, J. L. Canovas, J. L. Hernandez-Ramos, R. Torres Moreno, A. Skarmeta, Privacy-preserving solutions for blockchain: Review and challenges, IEEE Access 7 (2019) 164908–164940.
- [59] A. S. Alavizadeh, S. H. Erfani, M. Mirabi, A. Sahafi, An efficient distributed and secure algorithm for transaction confirmation in IOTA using cloud computing, The Journal of Supercomputing 80 (2) (2024) 1491–1521. doi:10.1007/s11227-023-05525-4.  
URL <https://link.springer.com/10.1007/s11227-023-05525-4>
- [60] D. Chaum, E. Van Heyst, Group signatures, in: Workshop on the Theory and Application of Cryptographic Techniques, Springer, 1991, pp. 257–265.
- [61] R. L. Rivest, A. Shamir, Y. Tauman, How to leak a secret, in: C. Boyd (Ed.), Advances in Cryptology — ASIACRYPT 2001, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 552–565.
- [62] S. Goldwasser, S. Micali, C. Rackoff, The knowledge complexity of interactive proof-systems (extended abstract), in: R. Sedgewick (Ed.), Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA, ACM, 1985, pp. 291–304. doi:10.1145/22145.22178.  
URL <https://doi.org/10.1145/22145.22178>

- 
- [63] J. Camenisch, M. Drijvers, A. Lehmann, Universally composable direct anonymous attestation, in: C. Cheng, K. Chung, G. Persiano, B. Yang (Eds.), Public-Key Cryptography - PKC 2016 - 19th IACR International Conference on Practice and Theory in Public-Key Cryptography, Taipei, Taiwan, March 6-9, 2016, Proceedings, Part II, Vol. 9615 of Lecture Notes in Computer Science, Springer, 2016, pp. 234–264.
- [64] H. K. Maji, M. Prabhakaran, M. Rosulek, Attribute-based signatures, in: A. Kiayias (Ed.), Topics in Cryptology – CT-RSA 2011, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 376–392.
- [65] C. Paquin, G. Zaverucha, U-prove cryptographic specification v1. 1 (2011).
- [66] J. Camenisch, S. Mödersheim, D. Sommer, A formal model of identity mixer, in: International Workshop on Formal Methods for Industrial Critical Systems, Springer Heidelberg, 2010, pp. 198–214.
- [67] A. de la Piedra, J. Hoepman, P. Vullers, Towards a full-featured implementation of attribute based credentials on smart cards, in: D. Gritzalis, A. Kiayias, I. G. Askoxylakis (Eds.), Cryptology and Network Security - 13th International Conference, CANS 2014, Heraklion, Crete, Greece, October 22-24, 2014. Proceedings, Vol. 8813 of Lecture Notes in Computer Science, Springer, 2014, pp. 270–289.
- [68] Z. Zhang, M. Król, A. Sonnino, L. Zhang, E. Rivière, El passo: Efficient and lightweight privacy-preserving single sign on., Proc. Priv. Enhancing Technol. 2021 (2) (2021) 70–87.
- [69] J. Camenisch, S. Krenn, A. Lehmann, G. L. Mikkelsen, G. Neven, M. Ø. Pedersen, Formal treatment of privacy-enhancing credential systems, in: O. Dunkelman, L. Keliher (Eds.), Selected Areas in Cryptography - SAC 2015 - 22nd International Conference, Sackville, NB, Canada, August 12-14, 2015, Revised Selected Papers, Vol. 9566 of Lecture Notes in Computer Science, Springer, 2015, pp. 3–24.
- [70] S. Brands, Rethinking public key infrastructures and digital certificates: building in privacy, Mit Press, 2000.
- [71] J. Camenisch, A. Lysyanskaya, An efficient system for non-transferable anonymous credentials with optional anonymity revocation, in: B. Pfitzmann (Ed.), Advances in Cryptology — EUROCRYPT 2001, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 93–118.
- [72] J. Camenisch, A. Lysyanskaya, Signature schemes and anonymous credentials from bilinear maps, in: M. Franklin (Ed.), Advances in Cryptology – CRYPTO 2004, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 56–72. doi:10.1007/978-3-540-28628-8\_4.
- [73] J. Camenisch, E. V. Herreweghen, Design and implementation of the *idemix* anonymous credential system, in: V. Atluri (Ed.), Proceedings of the 9th ACM Conference on Computer and Communications Security, CCS 2002, Washington, DC, USA, November 18-22, 2002, ACM, 2002, pp. 21–30.

- 
- [74] D. Pointcheval, O. Sanders, Short randomizable signatures, in: Proceedings of the RSA Conference on Topics in Cryptology - CT-RSA 2016 - Volume 9610, Springer, Berlin, Heidelberg, 2016, p. 111–126. doi:10.1007/978-3-319-29485-8\_7.
- [75] D. Pointcheval, O. Sanders, Reassessing security of randomizable signatures, in: Proceedings of the RSA Conference on Topics in Cryptology - CT-RSA 2018 - Volume 10808, Springer, Berlin, Heidelberg, 2018, p. 319–338.
- [76] D. Boneh, X. Boyen, H. Shacham, Short Group Signatures, in: D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, M. Franklin (Eds.), Advances in Cryptology – CRYPTO 2004, Vol. 3152, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 41–55, series Title: Lecture Notes in Computer Science. doi:10.1007/978-3-540-28628-8\_3. URL [http://link.springer.com/10.1007/978-3-540-28628-8\\_3](http://link.springer.com/10.1007/978-3-540-28628-8_3)
- [77] M. H. Au, W. Susilo, Y. Mu, Constant-Size Dynamic k-TAA, in: D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, R. De Prisco, M. Yung (Eds.), Security and Cryptography for Networks, Vol. 4116, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 111–125, series Title: Lecture Notes in Computer Science. doi:10.1007/11832072\_8. URL [http://link.springer.com/10.1007/11832072\\_8](http://link.springer.com/10.1007/11832072_8)
- [78] C. Héban, D. Pointcheval, Traceable constant-size multi-authority credentials, in: C. Galdi, S. Jarecki (Eds.), Security and Cryptography for Networks - 13th International Conference, SCN 2022, Amalfi, Italy, September 12-14, 2022, Proceedings, Vol. 13409 of Lecture Notes in Computer Science, Springer, 2022, pp. 411–434. doi:10.1007/978-3-031-14791-3\_18. URL [https://doi.org/10.1007/978-3-031-14791-3\\_18](https://doi.org/10.1007/978-3-031-14791-3_18)
- [79] O. Sanders, Efficient redactable signature and application to anonymous credentials, in: A. Kiayias, M. Kohlweiss, P. Wallden, V. Zikas (Eds.), Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part II, Vol. 12111 of Lecture Notes in Computer Science, Springer, 2020, pp. 628–656. doi:10.1007/978-3-030-45388-6\_22. URL [https://doi.org/10.1007/978-3-030-45388-6\\_22](https://doi.org/10.1007/978-3-030-45388-6_22)
- [80] U. Haböck, S. Krenn, Breaking and fixing anonymous credentials for the cloud, in: Y. Mu, R. H. Deng, X. Huang (Eds.), Cryptology and Network Security - 18th International Conference, CANS 2019, Fuzhou, China, October 25-27, 2019, Proceedings, Vol. 11829 of Lecture Notes in Computer Science, Springer, 2019, pp. 249–269.
- [81] T. Looker, V. Kalos, A. Whitehead, M. Lodder, The bbs signature scheme (2023). URL <https://identity.foundation/bbs-signature/draft-irtf-cfrg-bbs-signatures>

- 
- [82] S. Tessaro, C. Zhu, Revisiting BBS Signatures, in: C. Hazay, M. Stam (Eds.), *Advances in Cryptology – EUROCRYPT 2023*, Vol. 14008, Springer Nature Switzerland, Cham, 2023, pp. 691–721, series Title: *Lecture Notes in Computer Science*. doi:10.1007/978-3-031-30589-4\_24.  
URL [https://link.springer.com/10.1007/978-3-031-30589-4\\_24](https://link.springer.com/10.1007/978-3-031-30589-4_24)
- [83] G. Fuchsbauer, C. Hanser, D. Slamanig, Structure-preserving signatures on equivalence classes and constant-size anonymous credentials, *J. Cryptol.* 32 (2) (2019) 498–546.
- [84] D. Chaum, Blind signatures for untraceable payments, in: D. Chaum, R. L. Rivest, A. T. Sherman (Eds.), *CRYPTO’82*, Plenum Press, New York, USA, 1982, pp. 199–203.
- [85] G. Fuchsbauer, C. Hanser, D. Slamanig, Practical Round-Optimal Blind Signatures in the Standard Model, in: R. Gennaro, M. Robshaw (Eds.), *Advances in Cryptology – CRYPTO 2015*, Vol. 9216, Springer Berlin Heidelberg, Berlin, Heidelberg, 2015, pp. 233–253, series Title: *Lecture Notes in Computer Science*. doi:10.1007/978-3-662-48000-7\_12.  
URL [http://link.springer.com/10.1007/978-3-662-48000-7\\_12](http://link.springer.com/10.1007/978-3-662-48000-7_12)
- [86] U. Feige, A. Fiat, A. Shamir, Zero-knowledge proofs of identity, *Journal of Cryptology* 1 (2) (1988) 77 – 94.
- [87] C. Baum, T. Frederiksen, J. Hesse, A. Lehmann, A. Yanai, PESTO: Proactively Secure Distributed Single Sign-On, or How to Trust a Hacked Server, in: *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, IEEE, Genoa, Italy, 2020, pp. 587–606. doi:10.1109/EuroSP48549.2020.00044.  
URL <https://ieeexplore.ieee.org/document/9230400/>
- [88] M. Keller, MP-SPDZ: A versatile framework for multi-party computation, in: J. Ligatti, X. Ou, J. Katz, G. Vigna (Eds.), *CCS ’20: 2020 ACM SIGSAC Conference on Computer and Communications Security*, Virtual Event, USA, November 9-13, 2020, ACM, 2020, pp. 1575–1590. doi:10.1145/3372297.3417872.  
URL <https://doi.org/10.1145/3372297.3417872>
- [89] T. K. Frederiksen, J. Hesse, B. Poettering, P. Towa, Attribute-based single sign-on: Secure, private, and efficient, *Cryptology ePrint Archive* (2023).
- [90] A. Sonnino, M. Al-Bassam, S. Bano, S. Meiklejohn, G. Danezis, Coconut: Threshold issuance selective disclosure credentials with applications to distributed ledgers, in: *26th Annual Network and Distributed System Security Symposium, NDSS 2019*, San Diego, California, USA, February 24-27, 2019, The Internet Society, 2019, pp. 1–15.  
URL <https://www.ndss-symposium.org/ndss-paper/coconut-threshold-issuance-selective-disclosure-credentials-with-applications-to-distributed-ledgers/>

- [91] J. Doerner, Y. Kondi, E. Lee, A. Shelat, L. Tyner, Threshold BBS+ Signatures for Distributed Anonymous Credential Issuance, in: 2023 IEEE Symposium on Security and Privacy (SP), IEEE, San Francisco, CA, USA, 2023, pp. 773–789. doi:10.1109/SP46215.2023.10179470.  
URL <https://ieeexplore.ieee.org/document/10179470/>
- [92] R. Gennaro, S. Goldfeder, B. Ithurnburn, Fully distributed group signatures (2019).  
URL [https://www.orbs.com/wp-content/uploads/2019/04/Crypto\\_Group\\_signatures-2.pdf](https://www.orbs.com/wp-content/uploads/2019/04/Crypto_Group_signatures-2.pdf)
- [93] L. Harn, Group-oriented  $(t, n)$  threshold digital signature scheme and digital multisignature, IEE Proceedings - Computers and Digital Techniques 141 (5) (1994) 307–313.
- [94] W. Lueks, G. Alpár, J. Hoepman, P. Vullers, Fast revocation of attribute-based credentials for both users and verifiers, in: H. Federrath, D. Gollmann (Eds.), ICT Systems Security and Privacy Protection - 30th IFIP TC 11 International Conference, SEC 2015, Hamburg, Germany, May 26-28, 2015, Proceedings, Vol. 455 of IFIP Advances in Information and Communication Technology, Springer, 2015, pp. 463–478.
- [95] T. Nakanishi, N. Funabiki, Verifier-local revocation group signature schemes with backward unlinkability from bilinear maps, in: B. K. Roy (Ed.), Advances in Cryptology - ASIACRYPT 2005, 11th International Conference on the Theory and Application of Cryptology and Information Security, Chennai, India, December 4-8, 2005, Proceedings, Vol. 3788 of Lecture Notes in Computer Science, Springer, 2005, pp. 533–548.
- [96] K. Rannenberg, J. Camenisch, A. Sabouri (Eds.), Attribute-based Credentials for Trust: Identity in the Information Society, Springer, 2015.
- [97] J. Camenisch, A. Lysyanskaya, An efficient system for non-transferable anonymous credentials with optional anonymity revocation, in: B. Pfitzmann (Ed.), Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding, Vol. 2045 of Lecture Notes in Computer Science, Springer, 2001, pp. 93–118. doi:10.1007/3-540-44987-6\_7.  
URL [https://doi.org/10.1007/3-540-44987-6\\_7](https://doi.org/10.1007/3-540-44987-6_7)
- [98] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, G. Maxwell, Bulletproofs: Short proofs for confidential transactions and more, in: 2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA, IEEE Computer Society, 2018, pp. 315–334. doi:10.1109/SP.2018.00020.  
URL <https://doi.org/10.1109/SP.2018.00020>
- [99] H. Lipmaa, On diophantine complexity and statistical zero-knowledge arguments, in: C. Laih (Ed.), Advances in Cryptology - ASIACRYPT 2003, 9th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, November 30 - December 4, 2003, Proceedings, Vol. 2894 of Lecture Notes in Computer Science, Springer, 2003, pp. 398–415.

- [100] W. Mao, Guaranteed correct sharing of integer factorization with off-line shareholders, in: H. Imai, Y. Zheng (Eds.), *Public Key Cryptography, First International Workshop on Practice and Theory in Public Key Cryptography, PKC '98*, Pacifico Yokohama, Japan, February 5-6, 1998, Proceedings, Vol. 1431 of *Lecture Notes in Computer Science*, Springer, 1998, pp. 60–71.
- [101] F. Boudot, Efficient proofs that a committed number lies in an interval, in: B. Preneel (Ed.), *Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques*, Bruges, Belgium, May 14-18, 2000, Proceeding, Vol. 1807 of *Lecture Notes in Computer Science*, Springer, 2000, pp. 431–444.
- [102] D. Benarroch, M. Campanelli, D. Fiore, K. Gurkan, D. Kolonelos, Zero-knowledge proofs for set membership: Efficient, succinct, modular, in: N. Borisov, C. Díaz (Eds.), *Financial Cryptography and Data Security - 25th International Conference, FC 2021, Virtual Event, March 1-5, 2021, Revised Selected Papers, Part I*, Vol. 12674 of *Lecture Notes in Computer Science*, Springer, 2021, pp. 393–414. doi:10.1007/978-3-662-64322-8\_19.  
URL [https://doi.org/10.1007/978-3-662-64322-8\\_19](https://doi.org/10.1007/978-3-662-64322-8_19)
- [103] G. Arfaoui, J.-F. Lalande, J. Traoré, N. Desmoulins, P. Berthomé, S. Gharout, A practical set-membership proof for privacy-preserving nfc mobile ticketing., *Proc. Priv. Enhancing Technol.* 2015 (2) (2015) 25–45.
- [104] J. Blömer, J. Bobolz, D. Diemert, F. Eidens, Updatable anonymous credentials and applications to incentive systems, in: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS '19*, Association for Computing Machinery, New York, NY, USA, 2019, p. 1671–1685. doi:10.1145/3319535.3354223.  
URL <https://doi.org/10.1145/3319535.3354223>
- [105] E. C. Crites, A. Lysyanskaya, Delegatable anonymous credentials from mercurial signatures, in: M. Matsui (Ed.), *Topics in Cryptology - CT-RSA 2019 - The Cryptographers' Track at the RSA Conference 2019*, San Francisco, CA, USA, March 4-8, 2019, Proceedings, Vol. 11405 of *Lecture Notes in Computer Science*, Springer, 2019, pp. 535–555.
- [106] J. Blömer, J. Bobolz, Delegatable attribute-based anonymous credentials from dynamically malleable signatures, in: B. Preneel, F. Vercauteren (Eds.), *Applied Cryptography and Network Security - 16th International Conference, ACNS 2018*, Leuven, Belgium, July 2-4, 2018, Proceedings, Vol. 10892 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 221–239.
- [107] M. Chase, M. Kohlweiss, A. Lysyanskaya, S. Meiklejohn, Malleable signatures: New definitions and delegatable anonymous credentials, in: *IEEE 27th Computer Security Foundations Symposium, CSF 2014*, Vienna, Austria, 19-22 July, 2014, IEEE Computer Society, 2014, pp. 199–213.



- 
- [108] J. Bobolz, F. Eidens, S. Krenn, S. Ramacher, K. Samelin, Issuer-hiding attribute-based credentials, in: M. Conti, M. Stevens, S. Krenn (Eds.), *Cryptology and Network Security - 20th International Conference, CANS 2021*, Vienna, Austria, December 13-15, 2021, Proceedings, Vol. 13099 of Lecture Notes in Computer Science, Springer, 2021, pp. 158–178.
- [109] S. Krenn, T. Lorünser, A. Salzer, C. Striecks, Towards attribute-based credentials in the cloud, in: S. Capkun, S. S. M. Chow (Eds.), *Cryptology and Network Security - 16th International Conference, CANS 2017*, Hong Kong, China, November 30 - December 2, 2017, Revised Selected Papers, Vol. 11261 of Lecture Notes in Computer Science, Springer, 2017, pp. 179–202. doi:10.1007/978-3-030-02641-7\\_9.
- [110] S. Ringers, E. R. Verheul, J. Hoepman, An efficient self-blindable attribute-based credential scheme, in: A. Kiayias (Ed.), *Financial Cryptography and Data Security - 21st International Conference, FC 2017*, Sliema, Malta, April 3-7, 2017, Revised Selected Papers, Vol. 10322 of Lecture Notes in Computer Science, Springer, 2017, pp. 3–20.
- [111] T. Baignères, P. Bichsel, R. R. Enderlein, H. Knudsen, K. Damgård, J. Jensen, G. Neven, J. Nielsen, P. Paillier, M. Stausholm, D4.2 final reference implementation, ABC4Trust project deliverable (2014).  
URL <https://abc4trust.eu/index.php/pub/deliverables/208-d4-2-final-reference-implementation>
- [112] M. Rosenberg, J. White, C. Garman, I. Miers, zk-creds: Flexible Anonymous Credentials from zkSNARKs and Existing Identity Infrastructure, in: *2023 IEEE Symposium on Security and Privacy (SP)*, IEEE, San Francisco, CA, USA, 2023, pp. 790–808. doi:10.1109/SP46215.2023.10179430.  
URL <https://ieeexplore.ieee.org/document/10179430/>
- [113] F. Baldimtsi, J. Camenisch, L. Hanzlik, S. Krenn, A. Lehmann, G. Neven, Recovering lost device-bound credentials, in: T. Malkin, V. Kolesnikov, A. B. Lewko, M. Polychronakis (Eds.), *Applied Cryptography and Network Security - 13th International Conference, ACNS 2015*, New York, NY, USA, June 2-5, 2015, Revised Selected Papers, Vol. 9092 of Lecture Notes in Computer Science, Springer, 2015, pp. 307–327. doi:10.1007/978-3-319-28166-7\\_15.  
URL [https://doi.org/10.1007/978-3-319-28166-7\\_15](https://doi.org/10.1007/978-3-319-28166-7_15)
- [114] G. L. M. et al., Technical implementation and feasibility, in: *Attribute-based Credentials for Trust: Identity in the Information Society*, Springer, 2015, pp. 255–317.
- [115] B. Larsen, N. E. Kassem, T. Giannetsos, I. Krontiris, S. Vasileiadis, L. Chen, Achieving Higher Level of Assurance in Privacy Preserving Identity Wallets, online document <http://www.ioanniskrontiris.de/publications/2023937.pdf> (2023).
- [116] R. Impagliazzo, S. M. More, Anonymous credentials with biometrically-enforced non-transferability, in: *Proceedings of the 2003 ACM workshop on Privacy in*

- the electronic society, ACM, Washington, DC, 2003, pp. 60–71. doi:10.1145/1005140.1005150.  
URL <https://dl.acm.org/doi/10.1145/1005140.1005150>
- [117] M. Blanton, W. M. P. Hudelson, Biometric-Based Non-transferable Anonymous Credentials, in: D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, S. Qing, C. J. Mitchell, G. Wang (Eds.), Information and Communications Security, Vol. 5927, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 165–180, series Title: Lecture Notes in Computer Science. doi:10.1007/978-3-642-11145-7\_14.  
URL [http://link.springer.com/10.1007/978-3-642-11145-7\\_14](http://link.springer.com/10.1007/978-3-642-11145-7_14)
- [118] Y. Dodis, L. Reyzin, A. Smith, Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data, in: T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, O. Nierstrasz, C. Pandu Rangan, B. Steffen, D. Terzopoulos, D. Tygar, M. Y. Vardi, C. Cachin, J. L. Camenisch (Eds.), Advances in Cryptology - EUROCRYPT 2004, Vol. 3027, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 523–540, series Title: Lecture Notes in Computer Science. doi:10.1007/978-3-540-24676-3\_31.  
URL [http://link.springer.com/10.1007/978-3-540-24676-3\\_31](http://link.springer.com/10.1007/978-3-540-24676-3_31)
- [119] M. Blanton, M. Aliasgari, On the (non-)reusability of fuzzy sketches and extractors and security in the computational setting, in: Proceedings of the International Conference on Security and Cryptography, 2011, pp. 68–77.
- [120] D. Bissessar, C. Adams, D. Liu, Using biometric key commitments to prevent unauthorized lending of cryptographic credentials, in: 2014 Twelfth Annual International Conference on Privacy, Security and Trust, IEEE, Toronto, ON, Canada, 2014, pp. 75–83. doi:10.1109/PST.2014.6890926.  
URL <http://ieeexplore.ieee.org/document/6890926/>
- [121] N. D. Sarier, Comments on biometric-based non-transferable credentials and their application in blockchain-based identity management, Computers & Security 105 (2021) 102243.
- [122] Y. Wen, S. Liu, Robustly Reusable Fuzzy Extractor from Standard Assumptions, in: T. Peyrin, S. Galbraith (Eds.), Advances in Cryptology – ASIACRYPT 2018, Vol. 11274, Springer International Publishing, Cham, 2018, pp. 459–489, series Title: Lecture Notes in Computer Science. doi:10.1007/978-3-030-03332-3\_17.  
URL [https://link.springer.com/10.1007/978-3-030-03332-3\\_17](https://link.springer.com/10.1007/978-3-030-03332-3_17)
- [123] R. Canetti, B. Fuller, O. Paneth, L. Reyzin, A. D. Smith, Reusable fuzzy extractors for low-entropy distributions, J. Cryptol. 34 (1) (2021) 2.
- [124] Q. Alamélou, P. Berthier, C. Cachet, S. Cauchie, B. Fuller, P. Gaborit, S. Simhadri, Pseudoentropic isometries: A new framework for fuzzy extractor reusability, in:

- J. Kim, G. Ahn, S. Kim, Y. Kim, J. López, T. Kim (Eds.), Proceedings of the 2018 on Asia Conference on Computer and Communications Security, AsiaCCS 2018, Incheon, Republic of Korea, June 04-08, 2018, ACM, 2018, pp. 673–684. doi:10.1145/3196494.3196530.  
URL <https://doi.org/10.1145/3196494.3196530>
- [125] J. H. Cheon, J. Jeong, D. Kim, J. Lee, A reusable fuzzy extractor with practical storage size: Modifying canetti et al.’s construction, in: W. Susilo, G. Yang (Eds.), ACISP 2018, Vol. 10946 of LNCS, Springer, 2018, pp. 28–44.
- [126] K. Zhang, H. Cui, Y. Yu, Facial template protection via lattice-based fuzzy extractors, IACR Cryptol. ePrint Arch. (2021) 1559.
- [127] A. Arakala, J. Jeffers, K. J. Horadam, Fuzzy extractors for minutiae-based fingerprint authentication, in: S. Lee, S. Z. Li (Eds.), ICB 2007, Vol. 4642 of LNCS, Springer, 2007, pp. 760–769. doi:10.1007/978-3-540-74549-5\_80.  
URL [https://doi.org/10.1007/978-3-540-74549-5\\_80](https://doi.org/10.1007/978-3-540-74549-5_80)
- [128] J. Hesse, N. Singh, A. Sorniotti, How to bind anonymous credentials to humans, Cryptology ePrint Archive, Paper 2023/853, <https://eprint.iacr.org/2023/853> (2023).  
URL <https://eprint.iacr.org/2023/853>
- [129] C. Adams, Achieving non-transferability in credential systems using hidden biometrics, Secur. Commun. Networks 4 (2) (2011) 195–206. doi:10.1002/sec.136.  
URL <https://doi.org/10.1002/sec.136>
- [130] S. Krenn, J. Orlicky, D. Slamanig, T. Trpišovský, Ribac: Strengthening access control systems for pandemic risk reduction while preserving privacy, in: Proceedings of the 18th International Conference on Availability, Reliability and Security, ARES ’23, Association for Computing Machinery, New York, NY, USA, 2023, pp. 1–9. doi:10.1145/3600160.3605039.  
URL <https://doi.org/10.1145/3600160.3605039>
- [131] Globalsign iot identity platform, <https://www.globalsign.com/es/internet-of-things/iot-identity-platform>, accessed: 2023-06-30 (2023).
- [132] J. Höglund, M. Furuhed, S. Raza, Lightweight certificate revocation for low-power iot with end-to-end security, Journal of Information Security and Applications 73 (2023) 103424. doi:10.1016/j.jisa.2023.103424.  
URL <https://doi.org/10.1016/j.jisa.2023.103424>
- [133] G. Cooper, B. Behm, A. Chakraborty, H. Kommalapati, G. Mandyam, H. T. ARM, W. Bartsch, Fido device onboard specification 1.1 (2021).
- [134] S. R. Niya, B. Jeffrey, B. Stiller, Kyot: Self-sovereign iot identification with a physically unclonable function, in: H. Tan, L. Khoukhi, S. Oteafy (Eds.), 45th IEEE Conference on Local Computer Networks, LCN 2020, Sydney, Australia, November

- 16-19, 2020, IEEE, 2020, pp. 485–490. doi:10.1109/LCN48667.2020.9314816.  
URL <https://doi.org/10.1109/LCN48667.2020.9314816>
- [135] M. A. Qureshi, A. Munir, PUF-IPA: A puf-based identity preserving protocol for internet of things authentication, in: IEEE 17th Annual Consumer Communications & Networking Conference, CCNC 2020, Las Vegas, NV, USA, January 10-13, 2020, IEEE, 2020, pp. 1–7. doi:10.1109/CCNC46108.2020.9045264.  
URL <https://doi.org/10.1109/CCNC46108.2020.9045264>
- [136] Y. Kortensniemi, D. Lagutin, T. Elo, N. Fotiou, Improving the privacy of iot with decentralised identifiers (dids), *J. Comput. Networks Commun.* 2019 (2019) 8706760:1–8706760:10. doi:10.1155/2019/8706760.  
URL <https://doi.org/10.1155/2019/8706760>
- [137] L. Cocco, R. Tonelli, M. Marchesi, A system proposal for information management in building sector based on bim, ssi, iot and blockchain, *Future Internet* 14 (5) (2022) 140. doi:10.3390/fi14050140.  
URL <https://doi.org/10.3390/fi14050140>
- [138] M. A. Bouras, Q. Lu, S. Dhelim, H. Ning, A lightweight blockchain-based iot identity management approach, *Future Internet* 13 (2) (2021) 24. doi:10.3390/fi13020024.  
URL <https://doi.org/10.3390/fi13020024>
- [139] M. Lücking, C. Fries, R. Lamberti, W. Stork, Decentralized identity and trust management framework for internet of things, in: IEEE International Conference on Blockchain and Cryptocurrency, ICBC 2020, Toronto, ON, Canada, May 2-6, 2020, IEEE, 2020, pp. 1–9. doi:10.1109/ICBC48266.2020.9169411.  
URL <https://doi.org/10.1109/ICBC48266.2020.9169411>
- [140] S. de Diego, C. Regueiro, G. Maciá-Fernández, Enabling identity for the iot-as-a-service business model, *IEEE Access* 9 (2021) 159965–159975. doi:10.1109/ACCESS.2021.3131012.  
URL <https://doi.org/10.1109/ACCESS.2021.3131012>
- [141] S. Venkatraman, S. Parvin, Developing an iot identity management system using blockchain, *Syst.* 10 (2) (2022) 39. doi:10.3390/systems10020039.  
URL <https://doi.org/10.3390/systems10020039>
- [142] M. Popa, S. M. Stoklossa, S. Mazumdar, ChainDiscipline - Towards a Blockchain-IoT-Based Self-Sovereign Identity Management Framework, *IEEE Transactions on Services Computing* 16 (5) (2023) 3238–3251. doi:10.1109/TSC.2023.3279871.  
URL <https://ieeexplore.ieee.org/document/10135155/>
- [143] K. Zhang, C. K. M. Lee, Y. P. Tsang, Stateless Blockchain-Based Lightweight Identity Management Architecture for Industrial IoT Applications, *IEEE Transactions on Industrial Informatics* (2024) 1–12doi:10.1109/TII.2024.3367364.  
URL <https://ieeexplore.ieee.org/document/10468559/>

- [144] K. Gu, W. Zhang, S. Lim, P. K. Sharma, Z. Al-Makhadmeh, A. Tolba, Reusable mesh signature scheme for protecting identity privacy of iot devices, *Sensors* 20 (3) (2020) 758. doi:10.3390/s20030758.  
URL <https://doi.org/10.3390/s20030758>
- [145] X. Boyen, Mesh signatures: How to leak a secret with unwitting and unwilling participants, in: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Springer, 2007, pp. 210–227.
- [146] Y. Yang, H. Cai, Z. Wei, H. Lu, K. R. Choo, Towards lightweight anonymous entity authentication for iot applications, in: J. K. Liu, R. Steinfeld (Eds.), *Information Security and Privacy - 21st Australasian Conference, ACISP 2016, Melbourne, VIC, Australia, July 4-6, 2016, Proceedings, Part I, Vol. 9722 of Lecture Notes in Computer Science*, Springer, 2016, pp. 265–280. doi:10.1007/978-3-319-40253-6\_16.  
URL [https://doi.org/10.1007/978-3-319-40253-6\\_16](https://doi.org/10.1007/978-3-319-40253-6_16)
- [147] L. Nguyen, Accumulators from bilinear pairings and applications, in: A. Menezes (Ed.), *Topics in Cryptology - CT-RSA 2005, The Cryptographers' Track at the RSA Conference 2005, San Francisco, CA, USA, February 14-18, 2005, Proceedings, Vol. 3376 of Lecture Notes in Computer Science*, Springer, 2005, pp. 275–292.
- [148] H. Choudhury, Hashxor: A lightweight scheme for identity privacy of iot devices in 5g mobile network, *Comput. Networks* 186 (2021) 107753. doi:10.1016/j.comnet.2020.107753.  
URL <https://doi.org/10.1016/j.comnet.2020.107753>
- [149] M. Akil, L. Islami, S. Fischer-Hübner, L. A. Martucci, A. Zuccato, Privacy-preserving identifiers for iot: A systematic literature review, *IEEE Access* 8 (2020) 168470–168485. doi:10.1109/ACCESS.2020.3023659.  
URL <https://doi.org/10.1109/ACCESS.2020.3023659>
- [150] J. L. C. Sanchez, J. B. Bernabé, A. F. Skarmeta, Integration of anonymous credential systems in iot constrained environments, *IEEE Access* 6 (2018) 4767–4778. doi:10.1109/ACCESS.2017.2788464.  
URL <https://doi.org/10.1109/ACCESS.2017.2788464>
- [151] J. L. C. Sanchez, J. B. Bernabé, A. F. Skarmeta, Towards privacy preserving data provenance for the internet of things, in: *4th IEEE World Forum on Internet of Things, WF-IoT 2018, Singapore, February 5-8, 2018, IEEE, 2018, pp. 41–46*. doi:10.1109/WF-IoT.2018.8355229.  
URL <https://doi.org/10.1109/WF-IoT.2018.8355229>
- [152] A. Alcaide, E. Palomar, J. Montero-Castillo, A. Ribagorda, Anonymous authentication for privacy-preserving iot target-driven applications, *Comput. Secur.* 37 (2013) 111–123. doi:10.1016/j.cose.2013.05.007.  
URL <https://doi.org/10.1016/j.cose.2013.05.007>

- 
- [153] X. J. Lin, L. Sun, H. Qu, Insecurity of an anonymous authentication for privacy-preserving iot target-driven applications, *Comput. Secur.* 48 (2015) 142–149. doi:10.1016/j.cose.2014.08.002.  
URL <https://doi.org/10.1016/j.cose.2014.08.002>
- [154] R. Naisse, G. Baldini, G. Steri, Y. Miyake, S. Kiyomoto, A. R. Biswas, An agent-based framework for informed consent in the internet of things, in: 2nd IEEE World Forum on Internet of Things, WF-IoT 2015, Milan, Italy, December 14-16, 2015, IEEE Computer Society, 2015, pp. 789–794. doi:10.1109/WF-IoT.2015.7389154.  
URL <https://doi.org/10.1109/WF-IoT.2015.7389154>
- [155] R. T. Moreno, J. G. Rodríguez, J. B. Bernabé, A. F. Skarmeta, A trusted approach for decentralised and privacy-preserving identity management, *IEEE Access* 9 (2021) 105788–105804. doi:10.1109/ACCESS.2021.3099837.  
URL <https://doi.org/10.1109/ACCESS.2021.3099837>
- [156] V. Stafford, Zero trust architecture, NIST special publication 800 (2020) 207.
- [157] C. Baum, T. K. Frederiksen, J. Hesse, A. Lehmann, A. Yanai, Pesto: Proactively secure distributed single sign-on, or how to trust a hacked server, *Cryptology ePrint Archive*, Report 2019/1470, <https://eprint.iacr.org/2019/1470> (2019).
- [158] A. Guillevic, S. Masson, E. Thomé, Cocks–pinch curves of embedding degrees five to eight and optimal ate pairing computation, *Designs, Codes and Cryptography* (2020) 1–35.
- [159] R. Barbulescu, S. Duquesne, Updating key size estimations for pairings, *Journal of Cryptology* 32 (4) (2019) 1298 – 1336.
- [160] E. Barker, A. Roginsky, Transitions: Recommendation for transitioning the use of cryptographic algorithms and key lengths, NIST Special Publication 800 (2011) 131A.
- [161] J. Camenisch, R. Chaabouni, A. Shelat, Efficient protocols for set membership and range proofs, in: J. Pieprzyk (Ed.), *Advances in Cryptology - ASIACRYPT 2008*, 14th International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, Australia, December 7-11, 2008. Proceedings, Vol. 5350 of Lecture Notes in Computer Science, Springer, 2008, pp. 234–252. doi:10.1007/978-3-540-89255-7\_15.  
URL [https://doi.org/10.1007/978-3-540-89255-7\\_15](https://doi.org/10.1007/978-3-540-89255-7_15)
- [162] H. Lipmaa, Lecture 9: Secret sharing, threshold cryptography, mpc, T-79.159 Cryptography and Data Security, Helsinki University of Technology (2 2016).  
URL <http://www.tcs.hut.fi/Studies/T-79.159/2004/slides/L9.pdf>
- [163] J. Camenisch, M. Kohlweiss, C. Soriente, An accumulator based on bilinear maps and efficient revocation for anonymous credentials, in: S. Jarecki, G. Tsudik (Eds.), *Public Key Cryptography - PKC 2009*, 12th International Conference on Practice and

- Theory in Public Key Cryptography, Irvine, CA, USA, March 18-20, 2009. Proceedings, Vol. 5443 of Lecture Notes in Computer Science, Springer, 2009, pp. 481–500. doi:10.1007/978-3-642-00468-1\\_27.  
URL [https://doi.org/10.1007/978-3-642-00468-1\\_27](https://doi.org/10.1007/978-3-642-00468-1_27)
- [164] J. Camenisch, A. Lysyanskaya, Dynamic accumulators and application to efficient revocation of anonymous credentials, in: M. Yung (Ed.), Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 2002, Proceedings, Vol. 2442 of Lecture Notes in Computer Science, Springer, 2002, pp. 61–76. doi:10.1007/3-540-45708-9\\_5.  
URL [https://doi.org/10.1007/3-540-45708-9\\_5](https://doi.org/10.1007/3-540-45708-9_5)
- [165] J. B. Bernabé, J. G. Rodríguez, S. Krenn, V. Liagkou, A. F. Skarmeta, R. Torres, Privacy-preserving identity management and applications to academic degree verification, in: M. Friedewald, S. Krenn, I. Schiering, S. Schiffner (Eds.), Privacy and Identity Management. Between Data Protection and Security - 16th IFIP WG 9.2, 9.6/11.7, 11.6/SIG 9.2.2 International Summer School, Privacy and Identity 2021, Virtual Event, August 16-20, 2021, Revised Selected Papers, Vol. 644 of IFIP Advances in Information and Communication Technology, Springer, 2021, pp. 33–46.
- [166] M. J. Dworkin, Sp 800-38d. recommendation for block cipher modes of operation: Galois/counter mode (gcm) and gmac, National Institute of Standards & Technology, 2007.
- [167] S. F. Ali, M. A. Khan, A. S. Aslam, Fingerprint matching, spoof and liveness detection: classification and literature review, Frontiers of Computer Science 15 (1) (2021) 1–18.
- [168] R. Yavatkar, D. Pendarakis, R. Guerin, A framework for policy-based admission control, Tech. rep., RFC Editor (2000).
- [169] C. Herder, M. M. Yu, F. Koushanfar, S. Devadas, Physical unclonable functions and applications: A tutorial, Proc. IEEE 102 (8) (2014) 1126–1141. doi:10.1109/JPROC.2014.2320516.  
URL <https://doi.org/10.1109/JPROC.2014.2320516>
- [170] K. Krilakis, D3.3 design of physical unclonable functions for idm, Tech. rep., ERATOSTHENES project (2020).  
URL <https://eratosthenes-project.eu/wp-content/uploads/2023/01/ERATOSTHENES-D3.3-Design-of-Physical-Unclonable-Functions-for-IdM-v1.0-FINAL.pdf>
- [171] J. L. H. Ramos, S. N. Matheu, A. Feraudo, G. Baldini, J. B. Bernabé, P. Yadav, A. F. Skarmeta, P. Bellavista, Defining the behavior of iot devices through the MUD standard: Review, challenges, and research directions, IEEE Access 9 (2021) 126265–126285. doi:10.1109/ACCESS.2021.3111477.  
URL <https://doi.org/10.1109/ACCESS.2021.3111477>

- [172] J. L. H. Ramos, A. J. Jara, L. Marín, A. F. Skarmeta-Gómez, Dcapbac: embedding authorization logic into smart things through ECC optimizations, *International Journal of Computer Mathematics* 93 (2) (2016) 345–366. doi:10.1080/00207160.2014.915316.  
URL <https://doi.org/10.1080/00207160.2014.915316>
- [173] M. Salter, R. Housley, Suite b profile for transport layer security (tls), Tech. rep., RFC Editor (2012).
- [174] Rdf dataset canonicalization, <https://www.w3.org/community/reports/credentials/CG-FINAL-rdf-dataset-canonicalization-20221009/>, accessed: 2023-06-30 (2022).





# Publications

- [P1] R. Torres Moreno, J. Bernal Bernabe, J. García Rodríguez, T. Kasper Frederiksen, M. Stausholm, N. Martínez, E. Sakkopoulos, N. Ponte, A. Skarmeta, The OLYMPUS Architecture—Oblivious Identity Management for Private User-Friendly Services, *Sensors* 20 (3) (2020) 945. doi:10.3390/s20030945.  
URL <https://www.mdpi.com/1424-8220/20/3/945>
- [P2] R. T. Moreno, J. G. Rodriguez, C. T. Lopez, J. B. Bernabe, A. Skarmeta, OLYMPUS: A distributed privacy-preserving identity management system, in: 2020 Global Internet of Things Summit (GIoTS), IEEE, Dublin, Ireland, 2020, pp. 1–6. doi:10.1109/GIoTTS49054.2020.9119663.  
URL <https://ieeexplore.ieee.org/document/9119663/>
- [P3] S. Daoudagh, E. Marchetti, V. Savarino, J. B. Bernabe, J. García-Rodríguez, R. T. Moreno, J. A. Martinez, A. F. Skarmeta, Data Protection by Design in the Context of Smart Cities: A Consent and Access Control Proposal, *Sensors* 21 (21) (2021) 7154. doi:10.3390/s21217154.  
URL <https://www.mdpi.com/1424-8220/21/21/7154>
- [P4] J. García-Rodríguez, R. Torres Moreno, J. Bernal Bernabe, A. Skarmeta, Implementation and evaluation of a privacy-preserving distributed ABC scheme based on multi-signatures, *Journal of Information Security and Applications* 62 (2021) 102971. doi:10.1016/j.jisa.2021.102971.  
URL <https://linkinghub.elsevier.com/retrieve/pii/S2214212621001824>
- [P5] J. García-Rodríguez, R. Torres Moreno, J. Bernal Bernabé, A. Skarmeta, Towards a standardized model for privacy-preserving Verifiable Credentials, in: Proceedings of the 16th International Conference on Availability, Reliability and Security, ACM, Vienna Austria, 2021, pp. 1–6. doi:10.1145/3465481.3469204.  
URL <https://dl.acm.org/doi/10.1145/3465481.3469204>
- [P6] R. T. Moreno, J. Garcia-Rodriguez, J. B. Bernabe, A. Skarmeta, A Trusted Approach for Decentralised and Privacy-Preserving Identity Management, *IEEE Access* 9 (2021) 105788–105804. doi:10.1109/ACCESS.2021.3099837.  
URL <https://ieeexplore.ieee.org/document/9495805/>
- [P7] J. B. Bernabe, J. García-Rodríguez, S. Krenn, V. Liagkou, A. Skarmeta, R. Torres, Privacy-Preserving Identity Management and Applications to Academic Degree Verification, in: M. Friedewald, S. Krenn, I. Schiering, S. Schiffner (Eds.),

- Privacy and Identity Management. Between Data Protection and Security, Vol. 644, Springer International Publishing, Cham, 2022, pp. 33–46, series Title: IFIP Advances in Information and Communication Technology. doi:10.1007/978-3-030-99100-5\_4.  
URL [https://link.springer.com/10.1007/978-3-030-99100-5\\_4](https://link.springer.com/10.1007/978-3-030-99100-5_4)
- [P8] J. García-Rodríguez, D. Goodman, S. Krenn, V. Liagkou, R. T. Moreno, From Research to Privacy-Preserving Industry Applications: Workshop Summary, in: F. Bieker, J. Meyer, S. Pape, I. Schiering, A. Weich (Eds.), Privacy and Identity Management, Vol. 671, Springer Nature Switzerland, Cham, 2023, pp. 21–33, series Title: IFIP Advances in Information and Communication Technology. doi:10.1007/978-3-031-31971-6\_3.  
URL [https://link.springer.com/10.1007/978-3-031-31971-6\\_3](https://link.springer.com/10.1007/978-3-031-31971-6_3)
- [P9] K. Loupos, H. Niavis, F. Michalopoulos, G. Misiakoulis, A. Skarmeta, J. Garcia, A. Palomares, H. Song, R. Dautov, F. Giampaolo, R. Mancilla, F. Costantino, D. Van Landuyt, S. Michiels, S. More, C. Xenakis, M. Bampatsikos, I. Politis, K. Krilakis, S. Vavilis, An inclusive Lifecycle Approach for IoT Devices Trust and Identity Management, in: Proceedings of the 18th International Conference on Availability, Reliability and Security, ACM, Benevento Italy, 2023, pp. 1–6. doi:10.1145/3600160.3605083.  
URL <https://dl.acm.org/doi/10.1145/3600160.3605083>
- [P10] J. García-Rodríguez, A. Skarmeta, A privacy-preserving attribute-based framework for IoT identity lifecycle management, Computer Networks 236 (2023) 110039. doi:10.1016/j.comnet.2023.110039.  
URL <https://linkinghub.elsevier.com/retrieve/pii/S138912862300484X>
- [P11] J. García-Rodríguez, S. Krenn, D. Slamanig, To pass or not to pass: Privacy-preserving physical access control, Computers & Security 136 (2024) 103566. doi:10.1016/j.cose.2023.103566.  
URL <https://linkinghub.elsevier.com/retrieve/pii/S0167404823004765>
- [P12] J. García-Rodríguez, S. Krenn, J. Bernal Bernabe, A. Skarmeta, Beyond selective disclosure: Extending distributed p-ABC implementations by commit-and-prove techniques, Computer Networks 248 (2024) 110498. doi:10.1016/j.comnet.2024.110498.  
URL <https://linkinghub.elsevier.com/retrieve/pii/S138912862400330X>

# Appendix A

## Abbreviations

Abbreviation	Meaning
bb-ABC	Biometric-Bound Attribute-Based Credentials
CL	Caménisch-Lysyanskaya
DID	Decentralized Identifier
DLT	Distributed Ledger Technologies
dp-ABC	Distributed Privacy-preserving Attribute-Based Credentials
EBSI	European Blockchain Services Infrastructure
eIDAS	Electronic IDentification, Authentication and trust Services
EUDIW	European Digital Identity Wallet
FIDO	FIDO's Device Onboard Specification
FE	Fuzzy Extractor
GDPR	General Data Protection Regulation
IdM	Identity Management
IdP	Identity Provider
IoT	Internet Of Things
MUD	Manufacturer Usage Description
NIZK	Non-Interactive Zero-Knowledge
p-ABC	Privacy-preserving Attribute-Based Credentials
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PET	Privacy-Enhancing Technologies
PKI	Public Key Infrastructure
PS	Pointcheval-Sanders
PS-MS	Pointcheval-Sanders Multi-Signatures
PUF	Physical Unclonable Functions
SSI	Self-Sovereign Identity
SSO	Single Sign-On
TEE	Trusted Execution Environment
VC	Verifiable Credentials
VDR	Verifiable Data Registry
W3C	World Wide Web Consortium
ZK	Zero-Knowledge

Table A.1: Abbreviations