Engineering Ecosystems of Systems: UML Profile, Credential Design, and risk-balanced Cellular Access Control

David Bissessar

Thesis submitted to the University of Ottawa
in partial Fulfillment of the requirements for the
Doctorate in Computer Science

School of Electrical Engineering and Computer Science
Faculty of Engineering
University of Ottawa

Dedication

Maman,

Cette thèse t'est dédiée. Elle a été inspirée et complétée en grande partie grâce à toi et l'inspiration que tu me donnes…. Dédiée à ta joie de vivre, à ton courage et ta sagesse devant les complexités de la vie, à ta joie de vivre et ta capacité d'apprécier le moment présent et d'y trouver de la beauté, à ta présence et à chaque "goutte d'eau" que tu offres aux tiens. Merci Maman, pour qui tu es.

Merci de m'avoir donné la vie et de l'inspiration au cours des années.

Un chapitre se termine. Que réserve le prochain…

Acknowledgements

I would like to thank my supervisor, Dr Carlisle Adams. Dr Adams, thank you for your guidance, thoughts and friendship throughout. I am grateful for every step of the journey. I am amazed. I would also like to thank my jury and review panel, Dr. Daniel Amyot, Dr. William Anderson, Dr. Burak Kantarci and Dr. Anil Somayaji.

I would like to thank and acknowledge Defence Research and Development Canada's Centre for Security Science (DRDC CSS), the Canada Border Services Agency (CBSA) and the Department of Immigration, Refugees and Citizenship Canada (IRCC) for the opportunity, funding and forward thinking that helped me materialize some of these thoughts.

Abstract

This thesis proposes an Ecosystem perspective for the engineering of SoS and CPS and illustrates the impact of this perspective in three areas of contribution category First, from a conceptual and Systems Engineering perspective, a conceptual framework including the Ecosystems of System Unified Language Modeling (EoS-UML) profile, a set of Ecosystem Ensemble Diagrams, the Arms :Length Trust Model and the Cyber Physical Threat Model are provided. Second, having established this conceptual view of the ecosystem, we recognize unique role of the cryptographic credentials within it, towards enabling the ecosystem long-term value proposition and acting as a value transfer agent, implementing careful balance of properties meet stakeholder needs. Third, we propose that the ecosystem computers can be used as a distributed compute engine to run Collaborative Algorithms. To demonstrate, we define access control scheme, risk-balanced Cellular Access Control (rbCAC). The rbCAC algorithm defines access control within a cyber-physical environment in a manner which balances cost, risk, and net utility in a multi-authority setting. rbCAC is demonstrated it in an Air Travel and Border Services scenario. Other domains are also discussed included air traffic control threat prevention from drone identity attacks in protected airspaces.
These contributions offer significant material for future development, ongoing credential and ecosystem design, including dynamic perimeters and continuous-time sampling, intelligent and self optimizing ecosystems, runtime collaborative platform design contracts and constraints, and analysis of APT attacks to SCADA systems using ecosystem approaches.

Contents

Figures

Tables

Glossary

Anonymous Credential (AC)	An AbC Scheme defined by Camenisch and Lysyanskaya
Attribute-based Credentials (AbC)	A generic credential type which includes cryptographic credentials supporting extended privacy functionality. AC and DC are examples of AbC schemes.
Attestation	A certification made by an issuer regarding an attribute and its value as it pertaining to a subject. Attestations make up a credential and are consumed by Verifiers during the verification process, as part of vetting the service requests.
Attribute Value	Alpha-numeric values describing a named property about a Subject
Authority	A stakeholder within the ecosystem that may provide services acting in transactions, or may be an observer overseeing ecosystem activities. (Also called “Department” and “Service Provider” which may be used synonymously depending on the context.)
Blinding Protocol	An optional protocol conducted by the Holder of a credential whose goal is the obfuscation of the credential to prevent traceability between issuance and verification protocols. Blinding protocols are defined in AbC often in relation to the digital signature on the credential.
Component-Systems	Abstraction used in this thesis to describe the systems represented by the second “S” in the moniker “SoS”
Composite-System	Abstraction used in this thesis to help distinguish the differences between the two “Systems” in “SoS”. The composite-system is a synonym for the term “ecosystem” in EoS.
Collaborative Execution environment	A multi-stakeholder distributed computation model in which collaborative algorithms may be deployed over independent component systems and executed over the long term with objectively comparable performance toward expected which with measurable expectation toward local objectives to be made in a balanced manner(6.3.1).
Concourse	A logical construct within a CPE which projects the set of possible states that a transaction set may bring an ecosystem throughout its lifecycle. (1.6.1)
Cryptographic Credential (CC)	A collection of attribute attestations about a subject signed by an issuer, used for the asynchronous acquisition and redemption of reliable data for consideration by a verifier. Two main types are recognized in this thesis: EbC and AbC.

Cyber-Physical Ecosystem (CPE)	Within this thesis, a CPE is a multi-stakeholder SoS inscribed on an electronically connected physical terrain (3.3).
Cyber-Physical System	A CPS is a type of SoS implemented in a physical setting characterized by data connectivity and sensor-rich computer systems interacting with the physical terrain.
Cyber-Physical Social System	A CPS which recognizes and studies human and social aspects of CPS learning from humans-in-the-loop behavior of deployed sensors.
Cyber-Physical Social Services Ecosystem (CPSSE)	An EoS in which stakeholders transact predominantly for services rather than asset exchange. CPSSE is of interest in Government Services arena
Digital Credential (DC)	An AbC Scheme defined by Stefan Brands.
Ecosystem (See EoS)	In the context of this thesis, the word ecosystem is synonymous with “Ecosystem of Systems” which can be an IT Ecosystem or a cyber-physical ecosystem.
Ecosystem of Systems (EoS)	Technological distinction on SoS in which component systems are recognized to be subservient to the interests of the stakeholders who drive them.
Ecosystem of Systems - Unified Modeling Language Profile (EoS-UML)	A UML Profile which allows SoS to be modelled in terms of EoS concepts.
Envelope-based Credential	A traditional credential in which one signature certifies a collection of attributes. This is a common pattern is used in mobile Driver’s Licenses (mDL) and Digital Travel Credentials (DTC)
Issuer	An authority which creates verifiable attestations and credentials on subject attributes. Implements proofing, signature, and public key distribution processes.
Level of Assurance	A quality on an attestation which communicates to stakeholders of the ecosystem the measured confidence that proofed and attested value is correct. The ecosystem may rate a level of assurance. A transaction may also publish a level of assurance as a post condition.
Point of Presence	A connected compute device used by stakeholders to manage transactions within the cyber-physical ecosystem. The subject’s device and a service provider kiosk are examples of points of presence.
Privacy Attribute-based Credential (PABC)	See AbC.
Proofing	Out-of-band process conducted by a credential issuer to become convinced of the veracity of the attributes supporting a credential to be issued.

Service Request	A subject initiated request for service or consideration which is backed by a cryptographic credential. A service request on a credential involves some sort of verification protocol conducted by the service provider granting the request.
Subject	An entity described by a set of attributes, the subject of a credential. The subject is typically a participant in the issuance, storage, blinding, and redemption protocols. A stakeholder in a CPE. The subject may also be animate or inanimate.
Stakeholder	A human (or group of humans) who has an interest in the functioning of the ecosystem. Stakeholders may be subjects, authorities, observers, or designers within the ecosystem.
Storage Protocol	A protocol controlled by the credential holder which allows reliable and secure storage of an issued credential for later management and redemption in a service transaction.
Submission	Protocol by which the Holder submits a request for attestation to a provider of certified attributes
Transactor	Entities involved in a service request transaction.
Verification Protocol	A protocol that allows a service provider to verify the veracity of attestations submitted by a subject in support of a request for service or consideration. The verification protocol may include request preparation, blinding, and integrity check steps.
Verifier	The stakeholder which receives a cryptographic credential as part of a service request and must verify its integrity ownership and data compliance.

Notation

Element	Notation	Description
Countries		
	C_O	Country of Origin
	C_D	Country of Destination
Stakeholders		
	$s_i \in S$	A traveler, bearing passport of C_O
	a_0	$Governance$: C_O in alliance with C_D with representation of a_2
	a_1	Immigration Authority of C_D
	a_2	Airline preboarding verification at C_O
	a_3	Border Authority of C_D
	a_4	Law Enforcement of C_D
Contexts		
	z_1	Unconstrainted location within C_O
	z_2	Airport in C_O
	z_3	In-flight from C_O to C_D
	z_4	Unconstrainted location within C_B
Credentials		
	γ_0	e-Passport
	γ_1	the mobile Travel Authorization
Concourse Nodes		
	p_{11}	Travel Authorization issuance
	p_{21}	Pre-boarding main processing
	p_{22}	Pre-boarding overflow processing
	p_3	In-flight
	p_{41}	Customs Primary Inspection
	p_{42}	Customs Secondary Inspection
Credential Algorithms		
	Issue(…)	Issuance protocol between subject and issuer

	Blind(…)	Optional credential blinding protocol
	Verify(…)	Verification Protocol between subject and verifier
rbCAC		
	traverse(…)	Protocol allowing traversal of concourse
	assess_risk_category(…)	Policy specific risk assessment function
	next_state(…)	Risk based state transition
	policy_stabilize_risk(…)	Authority specific policy to balance cost and confidence of risk assessment of subject

Chapter 1 Introduction

The past 30 years have seen an evolution in networking technology for enterprises, governments, and individuals. In the early 1990s, enterprises focused on local area connectivity. The goals at the time were the sharing of resources such as printers and file servers, and the enabling of applications such as email and databases.

At the turn of the millennium, enterprises turned to internet technologies with web-delivery and thin-client approaches. Offerings such as dynamic content, e-commerce, extranets, and single sign-on were developed. Today, the focus is on reaching a larger client base, and a tighter integration of service-oriented architecture. Since about 2010, with the advent of distributed ledger technology, a new ecosystem approach has been in development. The focus has shifted from the intranet to the extranet to what might now be called the “Econet”.

We posit that an ecosystem perspective should be applied to the System of Systems (SoS) and its subtype, the Cyber-Physical Systems (CPS), and that in doing so, more robust, and stable SoS/CPS may be designed.

As a broad theme, this thesis proposes the concept of an “Ecosystem of Systems” (EoS) and maintains that the next important engineering refinements to enter the SoS/CPS world may be through what we call “ecosystem-thinking”, in which the EoS, like traditional software systems, has requirements and accountabilities to its

stakeholders, and that balancing the varied palette of user and stakeholder goals will be vital.

In the ecosystem, electronic transactions backed with cryptographic credentials (CC) help to provide this capability. The CC provides a way for issuers to make attestations about known subjects, for subjects to asynchronously accumulate and redeem these attestations with verifiers, who consume attestations in consideration of requested services.

We propose that the system of systems is better viewed as an ecosystem – thus becoming an Ecosystem of Systems (EoS). We demonstrate ecosystem-thinking by applying it to a Cyber-Physical System (CPS) to obtain what we refer to as a Cyber-Physical Ecosystem (CPE). The domain of focus for our CPE is a specialized area in Intelligent Transportation Systems (ITS) - the domain of Air Travel and Border Security (ATBS).

This thesis also proposes that EoS require a certification and transaction vehicle, much like currency, to mediate the trust issues between stakeholders. We propose the cryptographic credential as a mechanism to represent value, and facilitate transactions. We identify two broad types of cryptographic credential, the Envelope-based Credential (EbC) and the Attribute-based Credential (AbC), and demonstrate a credential design for each.

Finally, this thesis proposes that the compute resource of an EoS can be used as a distributed and collaborative environment, and we demonstrate this using a novel distributed access control algorithm named risk-balanced Cellular Access Control (rbCAC).

The current chapter sets the stage for the thesis by introducing the Cyber-Physical Ecosystem (CPE) and the sample domain of Air Travel and Border Security (ATBS), delineating the research problems, and introducing the focus areas of the credentials, target properties, threat model, and distributed processing environment.

1.1 Introducing the Ecosystem of Systems

The term “System of Systems” is commonplace and acceptable. Grammatically, it is a noun phrase which describes a composite-system (the first “S”) made up of distinct and independent component-systems (the second “S”, which is pluralized). In these SoS, the composite-system is the primary system of focus. The composite-system has particular properties: autonomy, belonging, connectivity, diversity, and emergence. Similarly, the component-systems have certain properties as well – they are self-sufficient systems, typically having independent owners, mandates, optimization, and maintenance schedules.

The term “System of Systems” is not without its problems. These have been noted in the literature by many, for example, by Maier (1997) and by Leveson (2013). Maier (1997) was among the first to point out problem in the term SoS, suggesting instead that the term “Collaborative Systems” might be better.

Maier’s notion of collaboration, which we embrace, is at the heart of the matter. However, the dynamics of “collaboration” are at one end of the spectrum in terms of productivity. Multi-stakeholder dynamics can also be adversarial or competitive, which are possibly counter-productive in an SoS but also, in general, an operational reality. This will be discussed in more detail in conjunction with our proposed Cyber-Physical Ecosystem Threat Model (CPE-TM). Moreover, in the two-word noun phrase

“collaborative systems” the notion of the composite-system (the first “S”) is obscured, implicitly referred to but with attention being removed from the system under study (the composite) and being drawn rather to the component-systems and a particular desired quality of their interactions.

Leveson (2013) elaborates on the problems of nomenclature, illustrating the dangers of viewing the composite-system differently than the component-systems, arguing that the composite system is equally accountable in terms of quality and safety as are the component-systems, and that naming them differently risks having negative impacts on the resultant quality and safety accountability of the composite.

The issues of quality and safety engineering as put forward by Leveson are central motivations of this thesis. We maintain however that the difference in the Systems referred to in the SoS moniker, is important enough that the quality and safety engineering within the parts and the whole benefit from having terms that distinguish them. For these reasons we propose a conceptual distinction that can be applied to SoS, viewing the composite-system, rather as an ecosystem. The composite system becomes an “ecosystem of systems” (or simply an “ecosystem” in this thesis).

This thesis defines an ecosystem-perspective for SOS – Ecosystem of Systems Unified Modeling Language (EoS UML). EoS UML adds a valuable perspective in engineering SoS. Each component-system in an SoS is independently administered and serves independent goals. Those goals are set by the stakeholder that owns them. Viewing the SoS as an ecosystem places the ecosystem designer-engineer in a role empathetic with those stakeholders. This allows the SoS to be designed such that it meets and is resilient to the different goals of the stakeholders. In doing so, the

composite-system has a better likelihood of being sustainable - its expected value over the long-term being stable enough for stakeholders to endorse with their participation.

EoS-UML and ecosystem-thinking is a set of concepts and techniques to be applied to the analysis and design of SoS/CPS to yield more robust resultant systems. These approaches do so by focusing discussion on the stakeholders and their objectives as cross-cutting characteristics to be used for ecosystem analysis and design and threat modeling (EAD&TM). EoS embraces the independence and social subservience of the component systems of an SoS with a view of better managing emergent behavior and complexity. EoS allows user-centric requirements analysis and iterative development and design methods to be applied to SoS and CPS development.

We illustrate the use of EoS-UML on a Cyber-Physical System in the Air Travel and Border Security (ATBS) domain. Applying ecosystem thinking and EoS-UML to CPS results in a more systematically constructed CPS – the Cyber-Physical Ecosystem.

1.2 Task Statements

The areas of focus in this thesis are as follows: 1) ecosystem conceptualization and threat model; 2) credential design and assessment; and 3) collaborative computing and algorithms (demonstration - access control). For each of these broad areas a number of specific research items are expressed in a problem-solving setting with salient questions raised, and scope of thesis discussed.

1.2.1 Task Statement 1: Cyber-Physical Ecosystem Modeling Statement

Task: Present a conceptual ecosystem model for multi-stakeholder CPS that can reflect contrasting hierarchical objectives and collaborative processing. Preferably, the model can be reused appropriately through the analysis, design, threat modelling and performance evaluation phases. Illustrate the model on the sample domain.

User Story: Clients frequently ask Daniel, the ecosystem designer, to design transaction and service models for their ecosystems. Initially, the business domains seem quite different. Over time, certain similarities emerge across domains. The similarities include remote attendance, the transactional nature of services, the need for confidentiality, privacy, and operational performance, and the conflicting goals of various stakeholders. **Research questions:** Can these cyber-physical ecosystems be approached in a general manner? What are the similarities that can be modeled? Moreover, do the threat models share any similarities?

1.2.2 Task Statement 2: Credential Design

Task: Describe envelope- and attribute-based design alternatives for the mobile Travel Authorization (mTA) in the ATBS domain. Select a collection of credential properties valued by the stakeholders of the selected ecosystem and evaluate the credential designs.

User Story: Alice s_1, a citizen of country C_o, seeks to visit the country of destination C_d for a vacation. According to the tourist travel requirements of C_d , visitors from C_o

must have a valid mTA, which can be obtained online, from the visa and immigration authority a_1 of C_d, using a mobile application downloaded to the traveler's smartphone. To create an mTA, the app captures e-passport data, a traveler selfie, and trip data. After these data are proofed by a_1, the mTA is issued and sent to s_1 for storage and later presentation to the airline a_2 (for verification at the time of pre-boarding) and by the border services agency a_3 (prior to entering C_d). The ATBS ecosystem is governed by a_0 a committee of delegates from C_O, C_d and a_1. The stakeholders s_i, a_1, a_2, a_3 have different credential requirements in terms of privacy, security, and operational qualities. A credential design is required which satisfactorily meets the requirements of the stakeholders.

Research questions: Can alternative credential design approaches be presented and analyzed against a set of illustrative properties? Are there trade-offs between designs? Is there a universally optimal design? Are there broad types of design that can be identified?

1.2.3 Task Statement 3: Risk-balanced Cellular Access Control

Task: Design a progressive access control system that operates in a SoS that is cost and risk sensitive. Provide a uniform interface such that organization-specific policies can be inserted. Capitalize on distributed computing resources and loosely coupled SoS. Assume a credential design with appropriate security, privacy, and operational properties.

User Story: Daniel has been commissioned by a_0 to design the ATBS mTA ecosystem. Having understood stakeholder goals and requirements for credential

representation, he examines the overall architecture and transaction processing model. Looking at the travel continuum, travellers obtain a visa, undergo screening by various authorities, and at any stage are permitted or denied access to their final goal. It seems the objectives of the governing group a_0 and the objectives of a_1 a_2 a_3 are designed to feasibly co-exist. The immigration authority's objectives are to maximize issued travel visas and minimize non-compliant travellers. The airline's objectives are to maximize flight occupancy minimizing disruptions. The border services objectives are to facilitate crossing for honest travellers and interdict threats. The country of destination wants to maximize tourism while maintaining public safety. It seems possible to hierarchically maximize global objectives while meeting the constraints and objectives of the local authorities. This could be done by optimizing the structure of the screening questions, wait times and risk reductions across the checkpoints in the continuum, organized in such a way that honest travelers enjoy easy, automated screening passage, whereas subject-travellers with a higher perceived risk would be more diligently screened, such that bad actors are interdicted early in the continuum. In such a system, traveller flow is specific to the perceived risk, thus each checkpoint guards its mandate, and hopefully can optimize performance targets. The overall ecosystem rules and thresholds, however, are calibrated such that the subject processing flow through the ecosystem converges toward a long-term equilibrium between, damage to the system, wait times, incorrect admission, and unacceptable inconvenience. If this could be achieved, the entire ecosystem cyber-physical resource would be acting in a collaborative manner to achieve an emerging global objective. Beyond ATBS, this

appears to be generally useful processing pattern. Consider credit applications with loan underwriting. A similar workflow of subject is processed by a multi-authority system, each of which performs a specialized task, which furthers the subject's path through the workflow. Each check, as well as the entire workflow has its objective function. A successful process and parameter configuration somehow balances between the various risks and utility of the system to the subjects and checkpoints.

Research questions: Can the progressive "permit", "detain" or "deny" decision with cost and risk sensitivity, be structured as an access control problem? Assuming authorities and the ecosystem owner can have conflicting goals, can an algorithm be defined that allows subject streaming to be conducted in a manner that achieves local objectives and optimizes overall ecosystem objectives? Can the access control problem be defined generically such that not only access control decisions can be performed? Can a processing model be created to treat the ecosystem cyber-physical sense-compute-actuate resource as a collaborative distributive processing engine?

1.3 Concept Background

In this section, general concepts related to CPEs and cryptographic credentials that are used throughout the thesis are introduced.

1.3.1 From "Systems of Systems" to "Ecosystems of Systems"

Figure 1 shows our target area of focus within the universe of systems. First off, the area of focus of the ecosystem thinking, broadly, is multi stakeholder systems. One

central premise to our thinking is that when multiple stakeholders are involved in a system's use, they must each have rational and justifiable expectation that a "satisficing[1]" proposition to their interests will be achieved. Beyond the multistakeholder dynamics of self-interest, this thesis focuses more precisely on Systems of Systems: systems which are independent systems from a mandate, managerial and maintenance perspective. From the SoS perspective, now, stakeholders are independent and have a compute-node presence for visibility into and participation in the composite system.

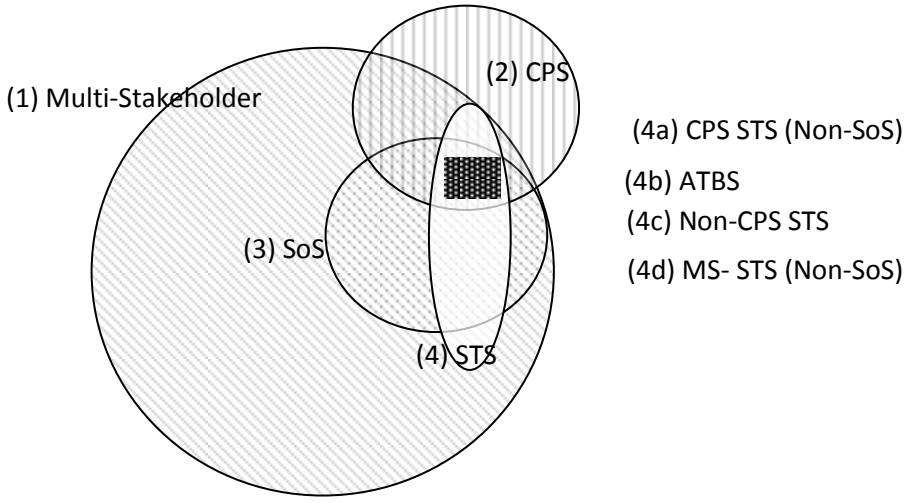

Figure 1) Universe of Systems and Area of Focus

The concepts in this thesis evolve from this starting point. Concepts of objective, transaction, risk-sensitive participation, and EoS-UML apply at this level. These come to

[1] The term "satisficing" was coined by Herbert Simon in "Administrative Behavior" REF _Ref69812393 \h . This term nicely crystalizes the concept of a balanced solution which attains long-term stakeholder satisfaction in some valid combination.

life when applied to the refined focus area at the intersection of SoS and CPS, namely, those systems that are composed of multistakeholder systems, and are inscribed on a geospatial terrain. A great variety of system themes are found in this area, for example national situational awareness and SCADA systems. Here, the concept of multiple stakeholders interacting in a geo-located cyber-physical arena come to life. Our proposal explicitly acknowledges the human user and consumer as intrinsic parts of the system, affecting its indeterminism and emergent behavior. The stakeholders are acknowledged and objectified within the dynamics of the composed system. In our proposal, we implicitly acknowledging the physical/virtual presence of the stakeholder, their interests, objectives, and intent in entering transactions, and the rational and intelligent self-determination of these stakeholders within the composed system over the long-term. As such, the proposal in this thesis is illustrated on ATBS, as an instance of a system of type (4b) from the Figure 1. When we discuss an “Ecosystem of Systems” we are precisely discussing systems at the intersection of Socio-Technical Systems (STS), Cyber-Physical Systems (CPS), Systems of Systems, and multistakeholder systems. We conjecture that the themes and techniques are of benefit, more largely, to the focus area, at the union of these system types than to the intersection, which is our focus.

1.3.2 Types of Cryptographic Credentials

Various types of credentials exist. In this section, the focus is on two alternative cryptographic credential designs, one in which an envelope-based approach is utilized and one in which an attribute-based approach is utilized.

1. **Envelope-Based Credentials**

Many credentials are defined by applying a standard digital signature scheme to a collection of multiple attributes. With envelope-based credentials, all the attributes certified by the issuer are placed, as a group, in an “envelope,” which is signed by the issuer, distributed to the holder, and later sent to the verifier as part of a third-party authentication or service request. The Digital Travel Credential (DTC) (ICAO 2020)of the International Civil Aviation Organization (ICAO) and the mobile Driver’s License (mDL) (ISO, 18013) are examples of signed, cleartext envelope-based credentials (EbC). Envelope-based credentials can be based on standardized signature schemes, which can ease widespread adoption.

2. Privacy Attribute-Based Credentials

Privacy attribute-based credentials (AbC) are a type of cryptographic credential data type in which specialized signatures are utilized to provide issuers, holders, and verifiers with functionalities such as selective disclosure, and composition of proofs using attributes from different credentials, non-traceable transactions, and zero-knowledge proofs of knowledge. Two of the most well-known AbC schemes in the field of cryptography are Stefan Brands’ digital credentials(Brands, 2000) , and Jan Camenisch and Anna Lysyanskaya’s anonymous credentials (Camenisch and Lysyanskaya, 2001). Digital credentials are at the basis of Microsoft’s UProve. Anonymous credentials form the basis of IBM’s IDEMIX and Evernym’s Sovrin.

3. Sign-on Credentials

The user-id/password combinations used for signing into enterprise and web systems are often referred to as “credentials” in the technology press. These “Sign-on Credentials” are a specific area of focus, with accompanying techniques, challenges,

and best practices. Related technologies include password-based access control, password salting, password strength rules, HTTP session propagation, single sign-on, and multifactor authentication.

The sign-on credential is generally a pair (pseudonym, secret). The password is usually stored in a back-end database. The envelope- and attribute-based credentials, in contrast, include multiple data elements, and request approval becomes a function of those attributes rather than a single secret stored in an authority's database. While there can be occasional overlaps in functionality and terminology, sign-on credentials are not the focus of this document.

1.3.3 Attribute-based Credential Functionality

1. **Signature randomization**

Signature randomization (or "blinding") is a AbC feature that provides a mechanism with which the holder can apply a signature to change its appearance so that it cannot be traced from the credential's issuance to its subsequent usage in service requests.

2. **Selective disclosure**

Selective disclosure is a specialized functionality whereby an issuer can certify and sign a collection of attributes, but the holder is free to divulge only the subset required by the verifier, while the signature still works. This is to be contrasted with the standard signature functionality noted above, in which all signed attributes must be disclosed for the signature to correctly verify.

3. **Zero-Knowledge Proofs of Knowledge**

A zero-knowledge proof of knowledge (ZKPOK) is a cryptographic protocol between a prover and verifier that enables the prover to convince a verifier of the truth of a

statement without disclosing anything about the attributes. For example, a prover might seek to prove a statement such as “The difference between my date of birth and today’s date is greater than eighteen years”, without divulging her actual birthdate. In the context of AbC, the ZKPOK also ties back to the signature of the issuer. Thus, the statement becomes: “My driver’s license bears a “date of birth” attribute which predates the year we are in now by at least 18 years. Furthermore, my driver’s license has a digital signature from the issuer (say, the ministry of transport) that attests to the veracity of the attribute.”

4. Credential Composition

Credential composition allows a holder to create a ZKPoK which combines selected attributes across credentials from different issuers (bearing different signatures). Combined data minimization and credential composition, together, permit “right-sizing” the dialog between verifier and the subject. The verifier now, may ask for only the attributes and assertions which are required, and the subject, in turn, need only reveal those attributes which are required. This is diametrically opposed to the same proof in an EbC scenario, in which the verifier and subject have no alternative but to provide all data in both credentials to obtain the subset of required information.

5. Combining the Functionality

Together, these functionalities provide the holder with powerful capabilities to express service requests, given a collection of credentials. The following examples demonstrate how a selective disclosure, credential composition, and zero knowledge proofs could be used: two credentials were issued to a citizen, an AbC driver’s license (“priv-DL”) and an AbC passport (“priv-PPT”).

Let us say, for example, that Alice wants to prove that she is a Canadian citizen who resides in the Ottawa area using her priv-DL and privacy-PPT. The statement to be proven might be:

> *"This photograph and name are mine. I am a Canadian citizen and I live in the Ottawa area. I have a valid privacy-PPT and priv-DL to support these statements."*

To properly convince the verifier, the proof disclosures and assertions might include some of the following:

a) *The photograph and name presented are from my privacy-PPT.*
b) *My name, as disclosed, is the same in my priv-DL and priv-PPT*
c) *In my priv-DL and priv-PPT, the date of birth attribute, though undisclosed, is identical.*
d) *In my priv-DL, the "city" field is either "Ottawa," "Gloucester," "Nepean," "Kanata," or "Orleans."*
e) *The credentials have not expired*
f) *The credentials have not been revoked*
g) *The priv-PPT is properly signed by the passport-issuing authority*
h) *The priv-DL is properly signed by the ministry of transport.*

In terms of credential features, the above demonstrates:

1. **Data minimization.** Alice divulges only her name, a photograph of her face, and the fact that she has a valid priv-PPT, a priv-DL, and a ZKPoK.
2. **Zero-Knowledge Proofs.** Alice proves, utilizing zero-knowledge, that the priv-DL and priv-PPT are valid, contain the divulged attributes, have equal valued attributes where necessary, and that she lives in the Ottawa area.
3. **Credential composition.** The attributes and ZKP produced by Alice span two credentials.

1.3.4 Credential Design Evaluation Properties

1.3.4.1 Conflicting Objectives

The stakeholders overseeing and transacting in an ecosystem have varied, possibly conflicting, objectives. Service providers, for example, might be interested in efficient processing, maximizing operating profits, and minimizing risks. Individuals, on the other hand, might be interested in increased privacy, efficiency, and convenience. The service provider's desire for low-risk transactions might lead to a data-hungry screening policy that is at odds with the individual's desire for an unobtrusive and privacy-respecting service. The choice of credential functionality will impact the various objectives of the stakeholders and the ecosystem in general. Attacks on those properties compromise the service levels of the ecosystem.

1.3.4.2 Target Properties

The ecosystem is engineered to deliver properties to the stakeholders. The design of the credentials is subservient to and enables these. For the sake of analysis and evaluation, many target properties are possible. Target properties change and evolve, limited only by the objectives of the ecosystem and its stakeholders, and the possibilities of technologies. Table 1 provides an illustrative set of target properties that are used throughout the thesis to demonstrate the impact of design decisions.

Table 1) Evaluation Properties

Security	Privacy	Operational
Unforgeability Tamper resistance Non-transferability	Unlinkability Composability Selective-show Biometric privacy	Interoperability Classifier accuracy Data usability Adaptability

1.4 An Instantiation of a CPE

This section comprises an introduction to the ATBS setting, which is used to illustrate the algorithms in this thesis.

1.4.1 Mobile Travel Credentials and Air Travel and Border Security

To introduce the mTA scenario, consider a traveler, who is equipped with a passport and a mobile phone, applying for a mobile credential from a destination country to visit as a tourist. The immigration authority of the destination country publishes an app that the traveler downloads and can use to obtain the mobile travel credential—a privacy-respecting cryptographic credential to be used when checking in at the airport to obtain a boarding pass, and again at the border to gain admission to the country of destination.

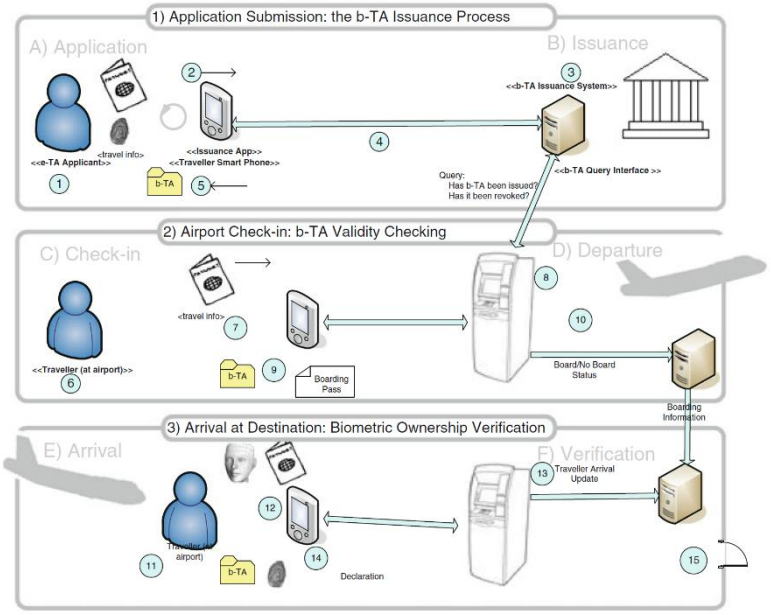

Figure 2) ATBS Checkpoint Workflow

Figure 2 demonstrates the checkpoint workflow in the travel continuum as consisting of three main steps from the traveler's point of view.

1) **Submission of Attributes and Issuance of mTA from a general access location**. The prospective traveler fills in an mTA application using her smartphone. The immigration authority receives and services the request remotely using an online system, which conducts risk screening and issues the required mTA.
2) **Presentation and Verification of mTA at the Airport of departure.** The airline verifies the traveler's documents when performing a pre-boarding risk assessment. The risk assessment might include biometric verification, document checks (lost or stolen passport verification), and board or no-board queries with the destination country.
3) **Presentation and Verification at the Airport of Arrival**. The border services authority verifies the identification, travel, and declaration documents at the border. A risk-sensitive decision is made regarding how the subject should be

processed. Should baggage be inspected? Should the subject be detained? Or should the subject be granted entry into the country of destination?

1.4.2 CPE Conceptual Model

In this thesis, ecosystems and business applications such as the ATBS and mTA are broken down into generic terms for technical analysis and design. For the purposes of this thesis, an ecosystem is a collection of regions of interest, each of which is controlled by an authority and conducts a screening of subjects who are seeking passage or a service. The subjects are possibly deceptive; thus, there is some inherent risk in erroneous *"permit"* or *"deny"* decisions. These decisions are subject to the Type 1 and Type 2 errors of classical statistics—the false accept and false reject. Differing authorities have differing operational tolerances for errors based on their risk-taking appetites.

Table 2) Mapping of ATBS to CPE Concepts

Generic CPE Concept	ATBS Concept
Authorities	Citizenship, immigration, airline, border services, destination country
Subjects	Honest and malicious travelers
Points of Presence	Smartphone application, Airline pre-board kiosk, Airline pre-board counter and terminal Border control kiosk Border control primary interview terminals Border control secondary interview terminals Supporting back-end systems
Credentials	e-Passport Mobile travel credential Airline ticket Boarding pass Declaration Entry stamp
Regions of Interest	Home country, airport of origin, pre-boarding area, airplane, airport of destination, border control area
Transactions	Traveler applying for mTA, Traveler requesting to board airplane Traveler requesting admission to destination country
False Accept	Allowing a traveler with a fraudulent identity or nefarious intention to obtain an

	mTA, board the airplane, or enter the country
False Reject	Preventing a well-intentioned traveler from obtaining an mTA, boarding the airplane, or entering the country

1.4.3 CPE Threat Model

The CPE threat model (CPE-TM) extends the standard threat model used in cryptography. The threat model mirrors characteristics of CPE, incorporating human, geospatial, and temporal elements. The proposed CPE attacker is composed of three components: command-and-control, a terrain attack force, and the platform. This approach to both the cyber-and physical-components of an attack, as well as the geo-located and temporal qualities of the algorithm execution and location platform (Chapter 7), are novel, and we believe, are important for CPS/SoS and distributed systems security in general. The CPE-TM can support modelling using scripts and graphical notation. This allows a staged attack to be tested prior to deployment. The effect of a cyber-physical attack might be stealthy, affecting either cyber or physical resources, and felt differently by various stakeholders within the ecosystem.

1.4.4 Credential Design Evaluation Properties

As noted above, multi-stakeholders may have conflicting objectives. Candidate credential and ecosystem designs will meet various requirements in different degrees. Attacks on those properties compromise the service levels of the ecosystem. For the sake of analysis and evaluation, many target properties are possible. Table 1 provides

an illustrative set of target properties that are used throughout the thesis to demonstrate the impact of design decisions.

1.4.4.1 Security Properties

Unforgeability: Unforgeability is the property that prevents the creation of a credential signature pair that bypasses the public key verification process by any party who is not in control of the issuer private signature key.

Tamper Resistance: Tamper resistance is the property that prevents a credential modification from bypassing the verification process undetected. Tamper resistance, for example, prevents an attacker from successfully submitting a credential in which the date of birth has been changed.

Non-transferability. Non-transferability (or non-lendability) is the property that prevents a credential issued to an individual from being used by anyone else.

1.4.4.2 Privacy Properties

Unlinkability: Unlinkability is the property that prevents an attacker from tracking the usage of the same credential of the same subject across multiple transactions while considering all available public data.

Biometric Privacy: The property of biometric privacy prevents the leakage of biometric information (templates, images, metric-space distances) to parties not involved in the transaction.

Selective Disclosure. The property of data selective disclosure permits the credential holder to divulge a subset of the attributes within a credential during the verification protocol. In this model, the holder controls the release of her

personal attributes. Verifiers are given only the attributes that are required for the transaction at hand.

Composability. The property of composability allows a subject to combine attributes from differing credentials in a verification proof. This property goes hand-in-hand with data minimization.

1.4.4.3 Operational Properties

Data Reliability. Authorities require assurances of data reliability to have confidence that the perceived risk and operational decisions that are made best suit reality.

Biometric Performance. The biometric performance of a system can be informally understood as its prediction accuracy in terms of its true or false, positive, or negative rate. The performance of biometric systems is an important consideration. A system should be able to be measured, configured, and deployed to predictably perform according to expected benchmarks.

Interoperability. An ecosystem serves multiple actors who might use differing algorithm implementations, depending on their operational needs. Thus, in an ecosystem, architectural, algorithmic, and interface choices that span system components are of paramount importance. The desire is to decouple components vs. constraining them. Architectural choices that decouple the separate component-systems, allowing them to select whichever algorithm they choose for face matching, for example, are preferable to architectural choices that force a wholesale adoption across the board. This is aligned with the SoS concept, in which the value of interoperability is emphasized.

Adaptability. The algorithms and components must be configurable by authorities and individuals to suit their needs. An algorithm choice at the ecosystem level that causes component systems to be brittle does not suit the security or operational needs of the stakeholders.

1.4.5 Assumptions

A number of assumptions are stated here. They are simply a starting point for the study. We will discuss relaxing some of these assumptions in Chapter 6.

[Assumption 1] Discrete-time sampling. The subject or environment is sampled at a discrete time interval. Continuous sampling can also be supported and is discussed in §7.3.2. However, its elaboration is beyond the scope of this thesis.

[Assumption 2] Fixed Contours. The contours of regions of interest are assumed to be fixed. While regions of interest might be chained and nested in complex manners, the contour of any one scope is assumed to be fixed through the lifecycle of the ecosystem. Dynamic or fluid contours are possible and are discussed in §7.3.2, below.

[Assumption 3] Conflict-free zone possession. The relationship between the stakeholder and zone is not in dispute throughout the ecosystem lifecycle.

[Assumption 4] Stable, Efficient Services. The services to regions of interest are stable. This includes communication, transcript, and directory services. Communications are assumed to be reliable. The transcript is assumed to provide security and privacy to the required specification. Thus, an attacker cannot target communications to temporarily deprive a zone or its

personnel of global public or obfuscated intelligence. Relaxing this assumption opens one up to a family of attacks, which, although interesting, are outside the scope of the current work.

[Assumption 5] Honest Ecosystem Designer. The ecosystem designer is assumed to be honest and to have the goal of establishing an ecosystem that has predictable states of expected benefit for all stakeholders. Periodic reinvestment is assumed to be a normal part of the lifecycle. Reinvestment budgets are assumed to be available when required. To use the ecosystem analogy, the Ecosystem Designer can be assumed an honest transactor in the Arms-Length Transaction Model in which the stakeholders are represented to the Ecosystem Designer as a requirements quorum. All required materials, and resource are assumed to be present, and the ecosystem designer is assumed to serve all stakeholders responsibly.

[Assumption 6] User Agents represent the Users. We assume that the computer platforms assigned to actors in the field represent their interests unless they have been corrupted. Thus, for example, a smartphone assigned to Alice does not seek to commoditize her data sharing or generate ad-revenue at the expense of her privacy or productivity. As another example, a handheld assigned to an officer in the field does not hang for vendor-defined updates, or prompt with fake news notifications. Per the CPE threat model, user agents might begin exhibiting this kind of behavior as they become corrupted. Chapter 4 departs from this assumption as it reflects work done on a government project, in which the

assumption of a secure and privacy respecting smartphone platform is not realistic.

[Assumption 7] Honest-but-Curious Service Providers.

In general, we assume that the verifier will faithfully participate in all protocols, etc., but it might also collude with other verifiers or with various authorities in order to learn information that it is not supposed to know. Specifically, with respect to biometric data, during the execution of a transaction, we assume that biometrics used for verification of identity are not stored beyond any stipulation in the protocol.

1.5 Discussion

While the ATBS ecosystem is the sample ecosystem in this thesis, the EoS-UML, the credential design approaches, the ecosystem as collaborative platform and rbCAC can be applied to other domains and application areas. This section comprises a brief discussion of some additional examples. Moreover, while this thesis has focus areas and contributions (model, credential design, and collaborative computing), there are a few broad themes that guide the overall work. These themes are briefly mentioned in this section.

1.5.1 Cyber-Physical Ecosystems in Other Domains

CPE can be applied in most multi-stakeholder ecosystem. The domains of virus control and remote piloted aircraft traffic management are excellent supplemental examples. Recall the features of the CPE: geo-physical terrain and rich sensor computation

devices (from CPS), human in the loop (from CPSS), multi-stakeholder or multi-objectives, and a transactional lifecycle with goals of sustainability.

In ATBS, the stakeholders are the airline, government agencies, and the traveler. A sample transaction lifecycle includes the mobile travel credential, in which certified data is tendered for services. The marketplace dynamics are characterized by service providers operating in a generally collaborative and synergistic manner, and a populace of subjects, the large majority of whom are well-behaved. We focus on discrete contexts and sampling time slices.

1.5.1.1 Virus Control

CPE and credentials can be applied in epidemiology—identifying incidence, distribution, and the control of disease in a population. The stakeholders include the public health authorities, policymakers, vaccine researchers, and those involved with virus isolation. The credential becomes a virus status credential borne by individuals and locations, alerting stakeholders to the virus status of the claimant. The terrain is the geographical area or city in which the controls are implemented.

The region of interest is the individual and social groups (e.g., the person, the family, and the educational cohort group), mobile inanimate units (e.g., the public bus or the airplane), fixed locations (e.g., the hotel room, the VIP floor, the residential building, or the restaurant), the civil administration unit (e.g., the community, the municipality, the province, and/or the country). The transaction lifecycle is the physical and cyber registration of individuals, with fine-grained location-based reporting during concourse navigation. Exceptional events are also geo-fenced and aggregative. The time-sampling domain is likely to be continuous in some applications and periodical in others, requiring

privacy and transparency, sometimes in a sensitive balance. Checkpoints then are centered on the transition points between them. This example illustrates the power of aggregative and recursive context definition, as well as strong pre-, post-, and invariant predicate clauses.

Successful attacks are numerous. Examples include the in-/ex-filtration of microbes across protected perimeters, the compromise of aggregate reporting totals, and the anonymity reversal of reporting entities.

1.5.1.2 Airspace Protection and Management.

For airspace control, the authorities include the civil aviation authority, the air navigation service providers, the airport authority, law enforcement, and safety council. The terrain includes land and air, and controlled and uncontrolled airspaces.

The regions of interest are the public and protected airspace and the sensitivity levels of the resources occupying the physical terrain below the airspace. The regions of interest demonstrate a recursive and aggregative composition. Computing and reporting distinctions regarding discrete vs. continuous time sampling and crispness of context boundaries, risk-level warnings, and dynamic intercession become increasingly important. Checkpoints are required to operate within perimeters with fluid boundaries. The credentials are numerous: they include aircraft, pilot and operator certificates, employee licenses, and service provider licenses. A successful attack on a component at a virtual perimeter can have a significant impact on operational performance, or cause harm to human life or facilities.

1.6 Contributions and Reading Guide

1.6.1 CPE Concept and Threat Model

We propose an extension of CPS in which human subjects conduct electronic transactions with multiple authorities in a cyber-physical ecosystem in which each stakeholder must weigh the potential costs and benefits of actions with respect to their own objectives. Towards formalizing this, we advance EoS-UML. To the best of our knowledge, no such formalism exists in the CPS literature.

We apply this formalism on a CPE of growing importance—that of air travel and border security. To accompany the CPE concept, we propose a threat model that, similarly, includes a cyber-physical attacker with human and machine elements with the ability to conduct distributed attacks through the CPS credential.

We advance a novel CPS threat model, in which the attacker has a threefold composition: a command-and-control attacker, a service platform of connected algorithms, and a distributed force of operatives on the field. This attacker is consistent with EoS-UML being objectively based and cyber-physical in nature. We refine previous definitions of the CPS threat model to account for stealth attackers, or attackers whose impact is not necessarily felt by the stakeholders of the system. To the best of our knowledge, our threat model is a novel approach. Its definition in UML and its application to the recursive definition of regions of interest (or contexts) facilitate the scripting of attacks (assuming parallel execution and reliable clock synchronization services), as well as pre-deployment testing of attacks in mock environments.

1.6.2 Credentials Design Alternatives and Analysis

We demonstrate the trade-offs that alternative credential designs might exhibit with respect to the goals of differing ecosystem stakeholders. We propose two designs, both of which achieve non-lendability through privacy-respecting biometric verification, with each having unique characteristics while considering security, privacy-sensitivity, and operational properties. One design has a fuzzy extractor secret as a biometrically derived key, which is then used to secure a symmetrically encrypted, asymmetrically signed credential. For the other design, we utilized derived key generation as an embedded secret in a Brands' digital credential scheme. The first design illustrates how the target mobile credential algorithms can be implemented in today's environment. The second design meets additional privacy by design requirements. These are both novel algorithms. Their application in the ATBS domain is also novel. In addition to these contributions, these algorithms are used to demonstrate the primary premise of this thesis, namely that, when engineering ecosystems, a fundamental consideration is the design of the digital credential vehicle that will be used to claim identity and privilege, and to conduct transactional risk analysis. In credential design, it is not the case that a given design will fit the desires of all stakeholders.

1.6.3 Risk-Balanced Cellular Access Control

We present a distributed, multi authority, risk-aware decision-making model and apply it to access control. Previous systems addressed risk-balanced access control in a single authority or single perimeter setting. In our setting, the perimeter is context-based, and the decision mechanism is dynamic. A subject is thus screened, not once, but

repeatedly throughout their interactions in the ecosystem. The progressive nature of our approach is also new. Since multiple stakeholders are involved in processing a subject, the work of one authority can benefit the screening processes of the next authority downstream. We demonstrate cellular access control in discrete time; however, it is also applicable in continuous time sampling. The result is a superior access control pattern to the traditional single-perimeter, single-authority approaches.

1.6.4 Thesis Structure

This thesis is structured as follows. Chapter 1 contains an outline of the contributions and conceptual framework. Chapter 2 comprises a description of the background of CPSSs, non-lendable digital credentials, and risk-sensitive distributed access control. Chapters 3 consists of a presentation of the UML metamodel, UML profile, ATBS application, target properties, and threat model. Chapters 4 and 5 comprise two alternative credential designs, one of which has an envelope-based pattern, and one that has an attribute-based pattern. These alternative designs are compared with respect to a set of security, privacy, and operational properties to highlight the differences and possible trade-offs. Chapter 6 comprises a presentation of the ecosystem as an architecture for risk-sensitive, collaborative, distributed processing, and demonstrates this capability on a proposed novel access control pattern, namely “risk-balanced cellular access control”. Chapter 7 is a discussion of CPE, comparative credential designs, and rbCAC . Chapter 8, which concludes the thesis, contains a description of ongoing and future work.

Chapter 2 Background

The thesis draws on several fields. Subsection 2.1 provides a background for our proposed Ecosystems of Systems (EoS) and Cyber-Physical Ecosystems (CPE). The subsection centers on System of Systems (SoS), Cyber-Physical Systems (CPS), Intelligent Transportation Systems (ITS), and Smart Cities. It introduces systems of systems (SoS), cyber-physical social systems (CPSS), and fog computing. Section 2.2 provides background information on the portrait-face modality of biometrics and classifier performance evaluation. Section 2.3 discusses Attribute-based Credential and fuzzy extractors. Section 2.4 presents background on risk-aware and multi-authority access control.

2.1 Systems of Systems and Cyber-Physical Systems

2.1.1 CPS and CPSS

(Sage and Cuppan, 2001) provide an early reference in the SoS literature. Although the study of CPS in air travel and border security is relatively new, there are commonalities shared between this field and Intelligent Transportation Services (ITS) and Supply Chain Management (SCM).

(Mathew, 2020) presents (ITS) as an SoS, a CPS, and a Cyber-Physical Social System (CPSS)(Xiong et al., 2015). Our target setting of ATBS is in line with the analysis of (Mathew, 2020) and exhibits similar properties, characteristics, and challenges. (Mathew, 2020) establishes that ITS clearly fall under the umbrella of CPS and CPSS (which are special types of SoS) and would benefit from research from this perspective.

(Xiong et al., 2015) point out the limitation of traditional CPS to quantitatively estimate and adapt to the impact of humans upon the system. They propose CPSS, which extend CPS by integrating social components. The model we present in this thesis further extends the CPSS perspective by considering multiple stakeholders (i.e., subject, authorities, and ecosystem owners) and their goals when faced with risk and uncertainty. CPE extends the “social” concept within CPSS by specifying that multiple stakeholders are necessarily involved, by including the possibility of conflicting objectives, as well as the ability for entities and components to appropriately collaborate when faced with risk. The complexity of analysis and modeling of SoS has been a widely acknowledged problem since the early work of (Jackson and Keys, 1984).Good surveys of social aspects of CPS exist (Zeng et al., 2020)(Anda and Amyot, 2019). From a CPS and design methodology perspective, (Zeng et al., 2020), identifies 1) that the design of CPSS is a complex task lacking in effective design approaches; and 2) that the security issues of CPSS have not been well investigated as they are still in their infancy. Our work contributes toward bridging these gaps. As Zeng observes, CPSS security investigation is at a nascent stage. In the area of CPS, several studies have been conducted. (Klötzer and Pflaum, 2015) present a research framework for CPS in logistics, Supply Chain Management Systems. They identify several characteristics of the solution presented in this chapter. These characteristics include the use of communications technologies such as Near-Field Communication (NFC), Radio-frequency Identification (RFID), mobile computing, and Cloud Computing, all of which enable the Internet of Things (IoT). This framework naturally applies to our CPE approach. Similar to CPSS, (Lenzini et al., 2015) also note the importance of

acknowledging the human in the loop and propose socio-technical physical systems. In their study they acknowledge the current state of need in CPS security analysis.

2.1.2 Fog Computing and IoT as Related to CPS.

The CPS body of knowledge has synergies with IoT and fog computing domains. CPS has traditionally been characterized as involving sensor-enabled devices with embedded software systems. Architectural patterns of fog, cloud and IoT also become pertinent as the devices gain power, greater network addressability, broader functionality, and a wider variety of edge-device form factors (such as medical devices, smartphones, and drones, for example).

A connected vehicle, a remotely piloted drone, or a smartphone bearing a food stamp wallet (all of which acquire, store, and dispense credentials) are all examples of edge or fog computing.

2.1.3 Cyber-Physical Security and Adaptive Multi-Stage Attacks.

(Nazarenko and Safdar, 2019) present a survey of privacy and cyber-security issues in CPS. This survey presents a thorough and systematic overview of the literature. A taxonomy of cyber-attacks is presented, which includes denial of service, eavesdropping, and malware. (Nazarenko and Safdar, 2019) also touch on physical issues of safety and on distribution and deployment concerns, reflecting the physical and geolocated nature of these SoS. The survey suggests that a compositional and encompassing model, such as we propose in the CPE attacker, is still not present in the literature.

Often the CPS literature originates from “traditional” cyber-attack taxonomies. (Loukas, 2015) presents a study of cyber physical attacks. Loukas distinguishes the cyber-physical attack from the well-studied cyber-attack in that a cyber-attack primarily affects confidentiality integrity and availability, whereas cyber-physical attacks are cyber-attacks which affect physical space by targeting elements such as the sensor, actuator, command/control, and other components of a cyber-physical system with purpose of physical effect such as disruption or damage of equipment or environment. Loukas lists broad categories of physical impact, including breach of privacy attacks (i.e., as would be caused by remotely hijacked sensors) and actuation attacks (i.e., unauthorized actuation, incorrect actuation, delayed actuation, and prevented actuation). Loukas also notes that most cyber-physical attacks include multiple breaches in cyberspace followed by multiple and possibly cascading effects in physical space.

This thesis finds the above definition of the cyber-physical attack to be unnecessarily limited, predominantly in that the effect is physical but also in the regards that the effect need not be felt by any or all the stakeholders alike, over any measurable timescales, etc. We propose to not limit the effect of the threat to the physical, or the time/sensor-perceptual.

Our CPE threat model re-emphasizes the dual nature, and in particular environmental attacks that can Affect digital perceptions on the field. Thus, the cross between cyber and physical such as changing illumination to alter biometric capture in a region of interest. Our cellular ecosystem model elaborates that the series of security

breaches may be achieved through a combination of cyber-hostile and physically hostile actions and that these may be composed often spanning security contexts.

Returning the focus to the perceptual. A particular kind of threat, the Advanced Persistent Threat (APT) operates on stealth, and is quite pertinent to our CPE threat. (Luh et al., 2013) present a survey of APTs. While many APTs are strictly cyber and do not have a physical impact, some certainly do. StuxNet, for example, targeted the proper functioning of uranium enrichment centrifuges in Iran.

(Hutchins et al. 2011) introduce APTs including the multi-stage kill chain, which is a stepwise generalization including steps for reconnaissance, weaponization, delivery, exploitation, installation, command and control, and actions on objective. (Cole, 2012) includes a kill-chain similar to that of (Hutchins et al. 2011), explicitly including the insider threat in the APT arsenal. The APT literature and our CPE model are synergistic. Some APT have traditional cyber objectives (i.e., corruption of data or theft of trade secrets). The CPE threat model we propose complements aspects of APTs such as multi-phased attack, goal orientation, the human element, time sensitivity, and possible stealth approaches. The multi-agent perspective of our CPE attacker allows the stealth aspects of APTs to be further dissected by encouraging the analyst to look for and consider the possibility of attack from non-obvious actors. The CPE attack as we propose, is expressive and scriptable. It can be structured as a series of steps analogous, but not limited to, the steps in a kill-chain.

(Huang and Zhu, 2018) examine adaptive defense strategies against APTs in cyber physical systems. They apply Bayesian games to model actions and counteractions in a multi-stage attack. Their model recognizes the uncertainty of the defender with respect

to the attacker's goals. This asymmetry of information is representative of a critical infrastructure defense position and a deceptive adversary. Their game theoretic approach is described as a two-player game over a multi-stage APT lifecycle over a finite time horizon. As examples, the authors refer to the multi-stage Stuxnet and Petya attacks. The two-player game by (Huang and Zhu, 2018) is synergistic to our CPE threat model and rbCAC (in which our Ecosystem Designer might be viewed as a single defender).

(Lykou et al., 2020) examines society and critical infrastructure for drone-related incidents near airspaces. The authors observe that countering drone attacks is a complex multi-step process involving the interaction between sensors, contexts, and human operators. However, the authors do not explore the multi-stakeholder nature of the airspace ecosystem. (Gorodetsky et al., 2008) explore the multi-agent nature of airspace incident management but do not examine threat modeling. Their model of New York airspace inspires a sample scenario in Chapter One.

2.1.4 Application Domains and ATBS

The air travel and border security (ATBS) domain is used as a backdrop to demonstrate concepts in this thesis. ATBS has similarities to intelligent transportation systems and supply chain systems, both of which are well-documented CPS. CPE includes characteristics and issues well-known to the CPS domain. (Lee, 2006) questions the effectiveness of current programming models in relation to the needs of CPS. These needs and challenges include time sensitivity in programming languages. These highlight problems between divisions of responsibility between operating

system and application, or software and hardware, as well as complexities in predictable memory allocation and hierarchy, pipelining and hyper threading, software modelling for concurrent components, reliable network timings, and blended approaches for system theories (including both hardware and software methodologies). Lee's challenges may be perceived as low-level system challenges, and that perhaps total ecosystem fitness requires more. It may be, that things like convergence analysis, equilibrium calibration, expected utility, context-sensitivity. adaptability, and reliability etc. all benefit from concepts at a higher level of abstraction. Concepts that encompass stakeholder goals such as “my operational mandate”, “my responsibility to my family”, “my quarterly results”, etc. This is where the thesis fits in. ATBS must contend with many of these considerations (for example, in relation to the reliability of time stamps on traveler records and the inherent concurrency of transaction and analytics processing by multiple authorities and travelers). In our architecture, these may be addressed by items such as ledger storage, transaction segregation across security domains, and coarse-grain separation between core service provider systems.

ATBS CPE builds on the components of today’s travel infrastructure; the e-passport, the verification kiosk, and integrated systems. Figure 3 presents an infographic by FRONTEX, Europe’s Coast Guard and Border Agency (FRONTEX, 2012). Ecosystem-perspective views this as a distributed, risk sensitive algorithm. The execution and rendering of the complete algorithm occur over time, distributed across component-systems running on the resources of a collection of component systems subservient to authorities with distinct, possibly conflicting objectives. The FRONTEX

diagram eloquently reflects the commonly adopted perspective that a travel and border screening decision is a multi-transactional decision, occurring over time and space. The traveller's intent to arrive at the border is communicated early, perhaps months before the actual date of arrival, as the traveler prepares and embarks on their trip. Current e-passport technology and the evolution of its security protocols as well as their capabilities and vulnerabilities in the areas of security and privacy are described in studies by (Bogari et al. 2012), (Hoepman et al., 2006), and (Juels et al., 2005). Security features include signatures on clear-text data stored on a chip in the passport against which kiosks can establish a secure channel for reading. Today's approach allows all information to be read by any reader that has been granted access.

Verification kiosk technology provides a multi-sensor way to authenticate a person using their e-passport and receive intent information such as a customs declaration. These kiosks and their capacities are discussed in (Nuppeney et al., 2010) and in work on AVATAR kiosk technology (Nunamaker et al., 2011) in which the kiosk becomes an adaptive multi-sensor interviewing station that changes its questioning based on the perceived intent of the subject. Currently projects in government and industry, such as Known Traveler Digital Identity (WEF, 2021) target distributed ledger credentials. The public literature of these projects attract attention, and foreshadow socio-technical appetites, however technical progress and approach are proprietary and not disclosed, neither peer-, public-, or cross-stakeholder reviewed. In general, the readiness of commercial and open-source offerings has not been assessed in terms of security, privacy, and operational requirements. As will become clear throughout this thesis, we

believe the assessment are an engineering necessity in the interest of the ecosystem and its stakeholders.

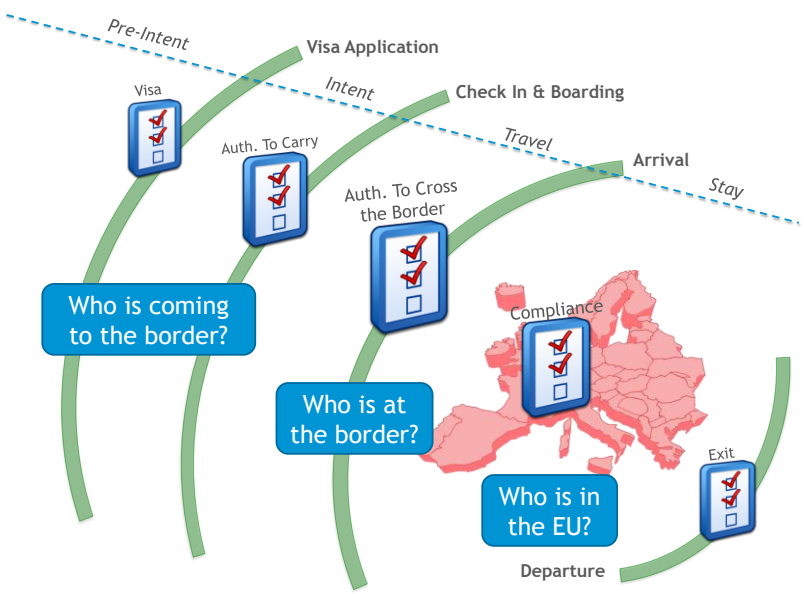

Figure 3) The Travel Continuum as a Cyber-Physical Social Ecosystem

CPE interposes structure and information processing from the earliest pre-departure phases of the travel continuum and allows a progressive update of risk assessment throughout. (FRONTEX, 2012)

Certain measures exist between airline and border controls to provide early information to enhance security and efficiency. One such measure is the advanced passenger information systems (APIS), which allows airlines to share information on incoming passengers to confirm boarding on scheduled aircrafts. There are differences in the way different countries implement their APIS interfaces; some countries provide a batch interface while others provide an interactive interface. In many cases, the countries have proprietary extensions from the standardized interfaces (WCO, 2013). These variances impede the efforts of airlines to conform

across the board. Our approach could alleviate such problems using a service bus approach with a common interface definition.

Similar to advanced passenger information (API) data is the passenger name record (PNR). PNR data sharing is subject to bilateral agreement between countries and has certain legal and privacy concerns (Banerjea-Brodeur, 2003)(Wilson, 2016). Our privacy by design solution means that consent-based sharing could alleviate some privacy issues. (Klötzer and Pflaum, 2015) observe key features in their analysis of supply chain management CPS. These include many of the features we observe in ATBS, including system components of time-dependent APIs, communication, automated kiosk screening, international-scale public key directories, and RFID chips with near-field communication (along with the estimated annual air traffic of six billion passengers).

2.2 Biometrics and Classifier Evaluation

The area of face biometrics and automated matching is a large field. Surveys have been conducted of algorithms and their strengths in different settings. This thesis is specifically concerned with a constrained subset of the problem space we refer to as passport-face biometrics. This subset of the larger area of facial biometrics can be considered as a biometric modality on its own. This section reviews previous work pertinent to the specific area of passport-face biometrics.

2.2.1 Biometric Verification

The field of biometrics is concerned with creating systems that can verify the match for a claimed identity (verification systems) or propose possible matches for an unknown identity (identification systems). This thesis focuses on nation-scale systems for 1:1 verification. Many countries have deployed biometric systems on a significant scale for security and citizen services. (Lehtonen and Aalto, 2017) discuss the deployment of automatic biometric verification kiosks in the EU and the views of stakeholders on such systems in terms of globalization, convenience, privacy, and security. Lehtonen and Aalto's work illustrates the trade-offs between stakeholder objectives that are central to our study.

Biometric systems can be presented as enablers of social improvement. India's Aadhaar is one such system. With an enrollment of over one billion citizens, Aadhaar is acknowledged as the world's largest biometric database. (Muralidharan, et al. 2016) discuss the size and benefit of India's biometric and universal identity program. They include discussions of how the program has facilitated such services as smartcard payments as well as citizen registration and biometric identification. Authors such as (Singh, 2019) and (Srinivas et al., 2020) highlight possible privacy concerns, social problems, and mission creep that may be associated with this program. Many nations are at various stages in their planning and deployment of various biometric identification and verification systems. (Khan, et al., 2010) present a study of biometrics and identity management in Saudi Arabia. The work of (Khan, et al., 2010) highlights the relationship between social services, homeland security, and identity management, of which biometric verification is a central feature. They also examine the dynamics this

may have on the citizen-government relationship. The authors review the various cards, usage scenario, and systems in place in the country and propose a streamlining approach. These tools and challenges are common internationally.

(Shaikh and Rabaiotti, 2010) discuss operational considerations in large scale biometric systems, identifying the properties of scale, accuracy, and privacy as operational characteristics in a trade-off relationship. The target properties against which our credential designs in Chapters Four and Five are evaluated include these as axes of comparison.

Passport-face biometrics. The specifications for the e-passport facial image are documented in a series of international standard documents. International Standards (ISO, 29794-5) specify standards for capturing, encoding, recording, transmitting, and quality of enrollment facial images used for e-passports.

Operational capture settings. While enrollment photos may undergo proofing and quality control procedures before they are finalized into the identity document, operational setting at the time of capture of the verification image can be harder to control. Specialized equipment may be configured in controlled areas, such as the automated border crossing (ABC) kiosk (FRONTEX, 2012). Using ABC technology, the subject can be provided with pose instructions. Specialized lighting and cameras may also be used. This thesis assumes such specialized equipment at the time of verification. The verification algorithms in Chapters Four and Five make use of kiosk technology at verification checkpoints. The rbCAC risk-sensitive interview can also be achieved with a kiosk checkpoint. In the case of rbCAC, presented in Chapter Six, an

AVATAR system such as that presented in (Nunamaker et al., 2011) may be used because it allows for risk-sensitive interviews and decision-making.

2.2.2 Focusing on the Frontal-Face Modality

The body of knowledge regarding face recognition is very large. Numerous surveys exist of the field (Zhao et al., 2003)(Li and Jain, 2011)(Sepas-Moghaddam et al., 2019)(Guo and Zhang, 2019). The focus of this thesis is a tightly constrained variation of the face recognition problem, which we refer to as the passport-face modality. In this passport-face modality, the typical setting is one of a single (unattended) enrollment image and a (generally also single) probe image, in which pose, expression, lighting, angle of capture, and distance between the eyes are specified by standards and best practice. These factors are controlled during the enrollment and verification process. This is not to suggest that passport-face verification algorithms are straightforward or error-free; however, the constrained modality helps to focus background discussion. Single-sample per person facial recognition is surveyed in (Tan et al., 2006).

2.2.3 Types of Face Recognition Algorithms

(Sepas-Moghaddam et al., 2019) present a multi-level taxonomy of face recognition algorithms which proposes a four-level approach to classifying face recognition algorithms: face structure representation (global, component plus face structure, component only), feature support (local or global), and feature extraction approach and sub-approach. Examining algorithms in terms of the third axis of categorization is useful in our summary. The categories based on feature extraction include appearance-based

algorithms, model-based approaches, learning-based approaches, and hand-crafted approaches. Examples of each of these are listed in Table 3.

Table 3: Categorization of Some Face Recognition Algorithms

Appearance-Based	Principal Component Analysis (PCA) Independent Component Analysis (ICA)	(Turk and Pentland, 1991b)
Model-based	Elastic Bunch Matching Graph (EBMG) 3D Morphable Model (3DMM)	(Wiskott et at., 1997) (Blanz and Vetter, 2003)
Learning-based	Deep Neural Nets Decision Pyramid (DP) Bayesian Patch Representation (BPR)	(Patkhi et al., 2015) (Zhang et al., 2017) (Li et al., 2016)
Hand-crafted	Local Shape Maps (LSM) Local Binary Patterns (LBP)	(Wu at al., 2004) (Ahonen et al., 2006)

Appearance-based techniques such as Principal Component Analysis (PCA) (Turk and Pentland, 1991b) and Independent Component Analysis (ICA) (Bartlett et al. 2002) map input images to a lower dimension representation. Appearance-based techniques are sensitive to specific qualities of the input images. As such, they are less tolerant of pose, expression, and lighting than other approaches. Model-based approaches such as the elastic bunch matching graph (EBGM) (Wiskott et at., 1997) and the 3D morphable model (3DMM) (Blanz and Vetter, 2003) seek to derive a geometrical representation of the face. Learning-based approaches derive data-driven relationships from training data sets; they tend to be more robust to variations such as scale, expression, and illumination. Custom approaches use pre-selected attributes, which are known to be relevant in the target domain, in a *"fit-for-purpose"* manner. Composite and hybrid approaches are also proposed. These include boosting, voting, and ensemble techniques (Viola and Jones, 2001)(Faltemier et al., 2008).

2.2.4 Eigenfaces as a Baseline Algorithm

The eigenfaces algorithm (Turk and Pentland, 1991a)(Turk and Pentland, 1991b) is a widely referenced baseline algorithm in face recognition. It is used in a number of fuzzy extractors proposals. Eigenfaces uses principal component analysis (PCA) on a gallery of like-dimensioned and pre-processed images to create a reduced-dimension set of basis vectors. Other facial images of the original dimensions can then be re-expressed as linear combinations of these lower dimension basis vectors. The coefficients of the linear combinations become the feature vectors. Biometric verification consists of comparing the Euclidian distance between feature vectors against a selected threshold.

While the technique is widely called eigenfaces (Turk and Pentland, 1991a), the technique of applying PCA to face images was first proposed by (Sirovich & Kirby, 1987). The presentation of the algorithm in (Sirovich & Kirby, 1987) is a complete and valuable reference.

The well-known work of (Turk and Pentland, 1991a) extends the work of (Sirovich & Kirby, 1987) by presenting a number of face processing considerations that have since become entire areas of research in themselves. These include algorithms for face detection, the discovery of unknown faces, the effects of lighting, pose, and expression, and techniques for real-time recognition. (Turk and Pentland, 1991a) is a seminal study in face recognition, touching on most of the complexities acknowledged in the field of face recognition today. The structure and content of (Li and Jain, 2011) exemplifies the relevance and depth of the areas highlighted in (Turk and Pentland, 1991a).

The Fuzzy Extractor construction of (Sutcu et al., 2007) can be used in the credential designs of Chapters Four and Five. The PCA algorithm is used to express a general biometric matcher in Chapters Four and Five.

2.2.5 Human Performance in Face Verification

(Philips and O'Toole, 2014) discuss differences between how humans and computers perform in face verification tasks. They propose a methodology and analyze a number of face recognition scenario and settings. When matching frontal faces in still images, algorithms are consistently superior to humans. For video and difficult still face pairs, humans are superior. In the testing protocol used, both machine and human subjects were presented with still images. However, this experiment is not illustrative of many operational situations. In the pre-boarding or customs situations discussed in this thesis, it is not the case that the operational personnel are presented with a still image at the time of verification. Instead, the identity claimant presents in person. Under these circumstances, we can capitalize on the superior ability of humans to consider extended information such as facial expression, skin texture, hair and 3D cranial structure and human behavior when making the identity verification/fraud detection decision.

It is precisely in recognition of these challenges that a staggered process is prototyped in Chapter Six. We propose a process in which both algorithmic and human matching are used in a synergistic manner. The practice of embedding human adjudication in data collection and decision processes for biometric evaluation is recognized as increasing system robustness (Grother et al., 2011).

2.2.6 Biometric Performance Assessment

The performance assessment of biometric systems is also a large area, drawing on the assessment of classifier systems (Japkowicz and Shah, 2011) and on techniques perfected in the domain of medical statistics (Altman and Bland 1994a)(Altman and Bland, 1994b)(Met, 1978). In this section we look specifically at techniques for the measurement and visualization of the performance of classifiers for binary decisions in 1:1 biometric verification in unbalanced datasets (Japkowicz and Stephen, 2002).

The Confusion Matrix

The confusion matrix captures correct and incorrect classifier decisions and allows the calculation of derived measures. Table 4 presents the basic measures, performance rates, test characteristics, and diagnostic ratios of the confusion matrix.

Table 4) Confusion Matrix Formulae

	Predictions +	Prediction -		
Condition +	True Positives (TP) "sensitivity"	False Negatives (FN)	True Positive Rate TPR = TP / Cond + "precision"	False Negative Rate FNR = FN / Cond +
Condition -	False Positives (FP)	True Negatives (TN) "specificity"	False Positive Rate FPR = FP / Cond -	True Negative Rate TNR = TN / Cond -
Prevalence Prev = Cond + / (Cond+ + Cond -)	Positive Predictive Value PPV = TP / Pred + "recall"	False Ommission Rate FOR = FN/Pred-	Positive Likelihood LR+ = TPR/FPR	Diagnostic Odds Ratio DOR = LR+/LR-
	False Discovery Rate FDR = FP/Pred+	Negative Predictive Value NPV = TN/Pred-	Negative Likelihood Ratio LR- = FNR/TNR	

At the heart of the confusion matrix are the number of correct and incorrect decisions made by the classifier. There are four such measures; true and false positives (TP and FP) and true and false negatives (TN, FN). As shown in Table 4, the rest of the measures can be derived from these.

ROC Curve Analysis

The Receiver-Operating Curve (ROC) curve (Altman and Bland 1994a)(Altman and Bland, 1994b)(Met, 1978)(Fawcett, 2006) provides a view of a decoder-classifier interpreting a received signal having possible distortion due to error. A ROC cure shows the rates of misclassification (the true positive rate (TPR) and false positive rate (FPR)) of a classifier, in terms of a population, at all classifier decision thresholds. Any point on along a ROC curve represents the underlying classifier's performance at a particular threshold. A point on the ROC curve corresponds to a complete confusion matrix. Since the performance is expressed in terms of rates, multiple classifiers can be compared on the same graph.

Figure 4 provides a sample ROC curve. Since the axes of the ROC curve are in rates, different classifiers can be plotted and compared on the same curve. A number of ROC curve features can be examined in order to understand classifier performance. The top left hand corner of the ROC curve represents ideal performance for a classifier (a false positive rate of zero and a true positive rate of 1). The point (0,1) is sometimes called ROC heaven. One way to judge a classifier is in terms of how closely it approaches this ideal point. This measurement can be given by the area under the curve (AUC). As a rule of thumb, the larger the AUC, the better the classifier. However, the area under the curve is not the decisive measure of classifier fitness. (Hand, 2009)

for example, suggests the AUC is almost meaningless. (Ferri et al., 2011) provides a rebuttal to (Hand, 2009) that includes discussions from the area of cost curves (Drummond and Holte, 2000).

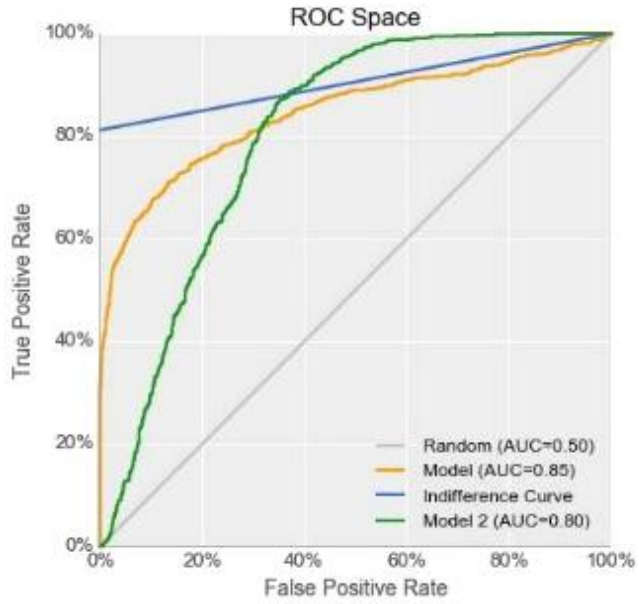

Figure 4) Comparative ROC Curve.

The ROC curve allows a manner to visualize of the performance of a classifier operating at different thresholds. Expressed in terms of rates, the ROC curve also allows classifier comparison.

Another way to use the ROC curve to infer classifier qualities to inspect the slope of the tangent at any given point. The tangent's slope indicates the expected change in FPR vs TPR for a threshold change in a given direction. The point at which the tangent is at 45 point is called the equal error rate (EER). According to the rule of thumb, the classifier with the EER closest to ROC heaven is the best. The problem with the EER is that it models a situation where organizational tolerance is equivalent for both false positives and false negatives. In practice, organizations will have their own indifference curve and it will rarely be EER.

Figure 4 highlights the interplay of TPR, FPR, AUC, and indifference curves. It shows the performance of two classifiers (represented by Model 1 in yellow and Model 2 in green), and an operational indifference curve (represented as a blue line with a slope of 0.2). Overall, Model 1 might seem best as it has an AUC of 0.85 (which is higher than the AUC of 0.80 achieved by Model 2).

Looking at the sample indifference curve, we see that its slope is less than 45 degrees. This means that the organization accepts a proportionally higher false positive rate for an increase in true positive rate. Model 1 approaches the indifference curve from below, whereas Model 2 exceeds it at times. Given the organization's goals, Model 2 is a superior classifier when it is configured to perform above the indifference curve.

There are number of alternatives to the ROC for visualizing classifier performance, include DET curves (Martin et al., 1997), cost curves (Drummond and Holte, 2000) and precision recall curves. The ROC curve will be used throughout this thesis as the baseline technology for verifiers to establish comparison thresholds that reflect the organization's tolerance for Type 1 and Type 2 errors. Nonetheless, we note that proper classifier and biometric performance assessment includes all considerations discussed in this section.

The issue of classifier performance evaluation is well-studied but is by no means a closed discussion. We note a discernable gap between standardization and evaluation of traditional biometrics and privacy-respecting biometrics. We are not aware of similar standard techniques in the privacy-respecting field. These have an impact on the fitness of fuzzy extractors for use in operations. This can be understood when considering

dynamic changes in match threshold, configurability, adaptability, and loose coupling between issuer and verifier systems.

The rbCAC algorithms proposed in Chapter Six can be seen as a classification algorithm; informally, the access control chain must properly classify subjects as “honest” or “suspicious” based on imperfect information when faced with deceptive behavior, and with costs when incorrect classifications are made. As a classifier, its performance can be evaluated using techniques from the above repertoire. We note that the selected domain and problem can be considered as a prediction of rare events problem that inherits the difficulties of unbalanced datasets.

A decision’s outcome can be seen as an interdiction of a threat or a facilitation of honest passage. As a facilitation, it consists of allowing honest claimants friction-free access; as an interdiction, it consists of stopping bad actors from disrupting the system. ATBS is similar to crime prediction, which is also characterized by unbalanced datasets where honest subjects constitute the overwhelming majority but the cost of a false positive may be quite high (Yu et al., 2011).

2.3 Privacy Credentials and Fuzzy Extractors

2.3.1 Privacy-Respecting Biometric Verification

Recent years have seen the growth of a new form of privacy-respecting biometric verification that does not require storage of the biometric or of a template derived from it. (Rathreb and Uhl, 2011) present a survey of this field. In this chapter, we present a mechanism that is designed to function with a secure sketch/fuzzy extractor for face biometrics such as that proposed by (Sutcu et al., 2009) or (Brien, 2020).

2.3.1.1 Fuzzy Extractors and Privacy-Respecting Biometric Verification

A number of surveys have been written in the field of privacy-respecting biometric authentication (Jain and Uludag 2003)(Uludag et al. 2004)(Jain et al. 2008)(Cavoukian and Stoianov 2009)(Rathreb and Uhl, 2011)(Sandhya and Prasad, 2017). There are a wide variety of algorithms in this area. Various techniques have been proposed that could be categorized in terms of the key generation approach, the method of addressing biometric variability (i.e., de-noising techniques or classifiers), or the distance metrics used. The methods proposed include the use of image transformations (Soutar et al., 1998)(Ratha et al., 2001), error correcting codes (Davida et al., 1998)(Juels and Wattenberg, 1999), quantization (Kevenaar et al., 2005), homomorphic encryption (Bringer and Chabanne, 2008), and fuzzy extractors (Juels and Wattenberg, 1999)(Juels and Sudan, 2002)(Dodis et al., 2004)(Rathgeb et al., 2013)(Rathgeb et al., 2014). In general, the goal is to perform biometric verification without storing or divulging the actual biometric. Work has been done toward the standardization of the terms used in privacy-respecting biometric verification schemes and the properties that these should exhibit (ISO, 24745)(Simoens et al., 2012). We limit our discussion to fuzzy extractors to support the credential designs in Chapters Four and Five.

In general, the fuzzy extractor allows the generation and regeneration of a key from two sufficiently close biometric samples. To help achieve this, the pair of functions defining the fuzzy extractor produce and use auxiliary “helper” data, (which is meant to be publicly storable), leaking no information about the biometric. (Davida et al., 1998) published the first paper in this area, using error correcting codes to introduce the concept of biometrically derived keys and to propose a verification architecture. (Juels

and Wattenberg, 1999) introduce the fuzzy commitment as a pair of functions to commit a secret and unlock it with a sufficiently similar version of the secret and stored public data. The hamming distance was used as the similarity metric. In follow-up, Juels and Sudan introduce the fuzzy vault, a construct of similar semantics designed to work using the set difference (Juels and Sudan, 2002). (Dodis et al., 2004) formalize fuzzy commitments and fuzzy vaults and define two primitives (the secure sketch and the fuzzy extractor). These are complimentary primitives that allow a strong key to be consistently regenerated from a noisy data stream.

The definitions of the secure sketches and fuzzy extractors are parameterized in terms of the error correcting capabilities of the underlying error correcting code, the security of the generated key, and the distance of the resulting public data from the random distribution. Security is stated in terms of min-entropy and conditional min-entropy, where min-entropy is a function of the probability of guessing the key and the conditional min-entropy is a function of the probability of guessing the key given the public data.

2.3.1.2 Concrete Fuzzy Extractor Schemes

Dodis' fuzzy extractor is a generic primitive. It has been applied to different biometric modalities and distance metrics, using different algorithms to resolve the noisy data. Fuzzy extractors have been designed for face, iris, fingerprint, and other biometric modalities. Noise resolution has been implemented using error correction codes, quantization, and classifiers. The proper choice of which fuzzy extractor to use depends on the template representation and distance metric for the selected biometric. Templates for iris biometrics are typically represented with discrete values, whereas

templates face images (as discussed for example with PCA) are represented with continuous values.

2.3.1.3 Templates on Discrete Values

Traditional fuzzy commitments and fuzzy extractors are defined to work on input vectors of discrete values. Many algorithms use biometric templates with continuous data. This is the case for the PCA algorithm described in Section 2.5. In (Li et al., 2006)(Sutcu et al., 2007) and (Sutcu et al., 2009), the authors propose the use of quantizers to convert biometric templates of real values into templates on integer values so that well-defined secure sketch schemes such as those defined by (Dodis et al., 2004) can be applied. In (Li et al., 2006) the authors illustrate an application of this technique using the frontal-face modality, singular vector decomposition (SVD) for template extraction, and the Essex 94 data set for testing. (Li et al., 2006) define a quantizer that uses a codebook which assumes the same error tolerances across all components and users. (Sutcu et al., 2007) generalize the construct to allow different error tolerances per component and per user. The authors also introduce a random projection approach to increase biometric performance.

(Sutcu et al., 2009) demonstrate the technique using PCA feature extraction (Sirovich & Kirby, 1987) and the ORL dataset (Samaria et al., 1994). In (Sutcu et al., 2009) the authors also further examine the effects of the random projections introduced in (Sutcu et al., 2007) on biometric performance and the cancelability of biometric keys. (Brien, 2020) examines fuzzy extractors for facial images, offers improvements on the construction of (Sutcu et al., 2007), and looks at using low density lattice codes (LDLC) as a quantizer in the fuzzy extractor (instead of using component- based quantization).

Chapters Four and Five assume a fuzzy extractor for faces such as that of (Sutcu et al., 2009)(Brien, 2020).

2.3.1.4 Related Approaches

This section introduces some biometric privacy techniques beyond the fuzzy extractor, which have influenced this thesis. (Adams, 2011) proposes a novel technique for privacy-respecting verification in which Pedersen commitments (Pedersen, 1992) are used on the bits of biometric templates and a series of zero-knowledge proofs of knowledge are used to determine acceptable Hamming distance of sets of commitments created at the time of enrollment and verification. In (Bissessar, 2013) and (Bissessar et al., 2014), the commitment-based approach of (Adams, 2011) is combined with biometric key generation. The key generated from a fuzzy extractor is stored in a Pederson commitment. At the time of verification, a fresh Pederson commitment is generated and key values are verified using a zero-knowledge proof.

2.3.1.5 Fuzzy Extractor Vulnerabilities and Limitations

The security and biometric performance of fuzzy extractors or particular constructs has been analyzed (Boyen, 2004)(Blanton and Aliasgari, 2011)(Liu et al., 2011)(Lafkih et al., 2015). (Boyen, 2004) identifies a multiplicity attack and proposes a countermeasure. In a multiplicity attack, an attacker can distinguish between users, and in some cases, reverse the biometric templates using the public data from multiple enrollments.

2.3.2 Credential Systems

Attribute-based Credential systems (AbCs or credential systems) provide protocols for individuals, issuers and verifiers to interact in the signing and verification of certified attributes. A number of systems enabling various nuances of privacy protection on identity, attributes, and transaction linkage have been proposed, including those of (Chaum, 1982)(Chaum, 1985)(Brands, 2000)(Camenisch and Lysyanskaya, 2001)(Verheuil, 2001). There are currently no systematic literature reviews of the field. However, (Koning et al., 2014) and (Veseli and Serna, 2016) provide an overview of the various functionalities available in AbCs and an appreciation of the performance of Microsoft's UProve and IBM's IDEMIX, which are based on (Brands, 2000) and (Camenisch and Lysyanskaya, 2001), respectively.

Digital credentials (DCs) are presented in (Brands, 2000). DCs are an attribute-based credential scheme, with security based on the discrete log and RSA problems. The protocol is simple and efficient, providing important features such as single-show unlinkability of transactions, selective show of attributes, composability, and unforgeability. Anonymous credentials (AC) are defined in (Camenisch and Lysyanskaya, 2001) and elaborated on in related papers (Camenisch and Lysyanskaya, 2002)(Camenisch and Lysyanskaya, 2004). Anonymous credentials are untraceable over issuance and multiple-show transactions. AC are based on the strong RSA problem and the decisional Diffie-Hellman problem.

Both schemes have noted shortcomings. ACs are a somewhat costlier protocol than DCs, requiring more exponentiations in both the issue and show protocols. The signature scheme in DC is based on the Fiat-Shamir heuristic (Fiat and Shamir, 1986).

A security weakness of the Fiat-Shamir heuristic is that no reduction to computationally hard problems is currently known. The current security gap of the Fiat-Shamir heuristic has been analyzed in (Goldwasser and Kalai, 2003)(Bitansky et al., 2013).

In terms of performance on resource constrained devices, DCs perform significantly better than ACs since they can be implemented on elliptic curves. (Bichsel et al., 2009) present an AC implementation, optimized for the Java card, in which an *"age-of-majority"* proof takes 16 seconds. While we do not know of a directly comparable empirical study for digital credentials, (Baldimtsi and Lysyanskaya, 2013b) gathers facts which provide a basis for estimation. Current RFID chip technology permits a 0.4 second elliptic curve multiplication, suggesting a significant improvement can be obtained using DC, compared to the 1.3 second performance numbers underlying the analysis in (Bichsel et al., 2009).

Non-transferability refers to the prevention of lending in DCs, which could otherwise simply be copied between persons. Several solutions to the lending problem have been proposed. These include disincentive measures such as the revocation of anonymity (Brands, 2000)(Camenisch and Lysyanskaya, 2001) or the embedding of valuable secrets such as a bank account number. Other solutions include preventative measures such as biometrics (Brands, 2000) and the use of privacy-respecting biometrics (Adams, 2011)(Bissessar, 2013)(Gerdes et al. 2016)(Sarier, 2021). Chapters Four and Five use privacy-respecting biometrics to prevent credential lending.

2.4 Risk-Balanced Cellular Access Control

Chapter Six presents a distributed cumulative form of risk-based access control. This subsection focuses on access control models that incorporate risk or that distribute the

decision over multiple parties. This is a relatively new field and our focus includes a unique combination of attributes.

2.4.1 Classical Access Control

While our setting is dynamic and multi-authority, many elements from classical access control can be used. For example, each authority may implement their own access control model, which may include mandatory aspects, a multi-layer approach, or a discretionary approach. Nonetheless, since our flow is characterized by a large number of travelers with a vast range of attributes and intents, many models will prove impractical. We thus propose a classifier-based method in which uncertainty of attributes and classification is a baseline to be acknowledged. While consistent with the access control models of (Lapadula and Bell, 1996) and Ferraiolo and Kuhn (1992) our approach is particularly well-suited to attribute-based access control (in which subjects and objects are associated with attributes and policies are defined in terms of these attributes).

The core concepts of the eXtensible Access Control Markup Language (XACML) (OASIS, 2013) are mirrored in our approach. As will be developed in Chapter Six, each context includes the architectural elements of a policy enforcement point, policy decision point, and hierarchical policies.

2.4.2 Risk-Aware Access Control

Much work exists on the subject of single-perimeter risk-aware access control, IoT single-perimeter access control, and collaborative game theoretic approaches.

However, we have found that very little work has been done at the intersection of these areas. The recent survey of (Atlam et al., 2020) on risk-based access control confirms our findings. The authors review 44 publications from an initial set of 1,044. Their findings seem to confirm a gap in research on decentralized Cellular Access Control. The survey provides the several takeaways. First, it highlights the distinction between traditional versus dynamic risk-based access control and the potential for flexibility and resilience. Second, it highlights the application of these methods in domains such as healthcare and the military (where thousands of lives can be at stake). Third, it highlights the differences between traditional and dynamic systems (with dynamic systems notably including context-based information for decision-making). Fourth, it confirms that many risk factors and risk estimation approaches exist and that these vary significantly based on the context or domain.

A number of approaches have been proposed to incorporate risk-based decisions into access control, including risk-adaptive access control (RAdAC) (McGraw, 2009)(Kandala and Sandhu, 2011) and risk-aware access control (RAAC). (Molloy et al., 2012) propose a single Policy Enforcement Point (PEP) (non-collaborative) mode which features a local decision point with a machine learning classifier, a measured level of uncertainty, and a methodology to derive a crisp action using an uncertainty resolution mechanism. This includes an oracle-like central policy decision point, the invoking of which includes a cost but dispels uncertainty.

2.4.3 Risk Estimation Techniques

The subject of risk estimation techniques in risk-based access control is examined by (Atlam et al., 2017). The selection found in the literature includes fuzzy logic, game theory, decision trees, Monte Carlo simulation, expert judgment, and formulaic approaches. Our approach is to delegate the choice of risk assessment method to the specific authority by including both mathematical models and expert judgment in the suite of possible techniques. This provides the human element that makes CPSS and CPE increasingly meaningful and non-deterministic. It is interesting to note that the information gain approach to attribute valuation and selection is a technique specifically from the field of decision trees.

2.4.4 Risk-Based Approaches in Business Ecosystems

Outside the technical literature, both risk-based and staged approaches are recognized in business ecosystems. Examples of these are plentiful, with insurance and the financial industry serving as obvious examples. In ATBS, as discussed by (Poole, 2008), risk is also recognized and incorporated into the functioning of the business ecosystem. (Degenhardt and Bourne, 2020) present the interplay of experience and science in developing border control technology for use in the front lines.

Chapter Six combines risk-aware methods of access control as surveyed by(Atlam et al., 2020) with the multi-actor game theoretic models of (Manshaei et al., 2013).

2.5 Chapter Summary

This chapter provides background information for the areas of CPS, biometrics and classifier evaluation, non-lendable digital credentials and privacy-respecting biometric verification, and risk-aware and multi-authority access control.

Cyber-physical systems are a type of SoS. The latter present widely acknowledged engineering challenges, which include emergent behavior and non-determinism. A distinction known as CPSS has recently arisen in the literature within CPS. CPSS recognize human in the loop as an important factor, which adds complexity to that already recognized in CPS. We further add the distinction of CPE to CPSS. While the addition may seem to complicate matters, part of our hypothesis is that recognizing the multi-party goal centricity of the usage of CPS assists in taking in to account the varied interests that are served by the well-designed ecosystem. The properties of security, privacy, and operational soundness, along with the characteristic of emergent behavior are focal points of this thesis.

Toward soundness of properties, we draw on the background of AbC and fuzzy extractors. The biometric modality of passport-face biometrics is widely deployed in ATBS and is therefore our focus. Biometric evaluation is not standardized in fuzzy extractors. PCA represents a useful baseline algorithm for face similarity matching because it is widely known, is applicable to passport-face biometrics, and is used at the core of some fuzzy extractors for face. In biometrics and classifiers, however, the field is well-studied and non-trivial. ROC analysis is identified as a best practice which can be used as a baseline for biometric and classifier evaluation.

The air travel and border security setting demonstrates a type of intelligent transportation system (ITS) which is a CPSS with multiple stakeholders and goal-driven behavior.

Several characteristics confirm that further study is merited, including the human aspects which add complexity in terms of non-deterministic behavior. The various goals of the system open the discussion for operations research techniques in optimization (subject to constraints). This motivates the work in Chapter Six on rbCAC. We do not implement optimization which may be achieved with simulated annealing, genetic programming. Goal oriented requirement languages (Anda and Amyot, 2019) may also be useful in integration with our techniques for hierarchical multi-objective optimization. .

Privacy-respecting designs for non-lendable credentials remains an open and interesting problem given the variety of business rules that can exist, and trade-offs present between stakeholders consuming the credentials. This breadth of possible approaches here may be analogous to that in the domain of digital signatures. The solutions in this thesis and the comparative approach to properties may further the study of these primitives.

Much of the literature on access control has been focused on single authorities. Not much work has been done on multiple authority or staged decision-making processes. Most risk-aware models are single-perimeter focused. To our knowledge, no work has proposed a decentralized model such as the rbCAC proposal in Chapter Six.

The following chapters elaborate specific designs for non-lendable digital credentials (Chapters Four and Five) and for the multi-authority distributed setting including risk-assessment and risk-balanced Cellular Access Control (Ch. 6).

Chapter 3 Ecosystem Ensemble Diagrams for SoS

3.1 Chapter Introduction

The Unified Modeling Language (UML) represents a general manner for modeling software systems. This chapter describes a UML metamodel and profile for multi-stakeholder ecosystems and then uses the profile to define a set of ecosystem domain and threat analysis diagrams, demonstrating their use on ATBS. Since the model and diagrams are explained, the roles and responsibilities of the entities that comprise an ecosystem are further elaborated, along with the assumed features of the ATBS sample. The concepts in this chapter set the stage for the subsequent Credential Design and Collaborative Processing chapters.

3.1.1 Contributions

3.1.1.1 Scope

1. EoS UML

We provide a profile EoS-UML with which CPS instance diagrams can be described using general high level terminology. Many instances of UML in the CPS literature are at a lower level, with focus on sensors, controllers and robotics (Magureanu et al., 2010)(Bagnato et al., 2017). This type of model is presented to help facilitate the discussion of the system after the fact. EoS-UML, in contrast, allows a top-down approach for the development of these systems. SysML provides a rich UML profile for

modelling sensors and activations and considering timing issues. Similarly, MARTRE provides a UML Profile for embedded CPS applications(Mallet et al., 2017), providing low-level modelling of meter, measure, and device with a case study on the embedded logic of a quadcopter. These frameworks operate at a different level of abstraction and component granularity than EoS-UML. EoS-UML focuses on coarse grained components and their transactional interactions and the goals of the stakeholders. EoS-UML complements offerings like SysML and MARTRE with ecosystem, transactional, and objective semantics.

.

2) The Ecosystem Ensemble Diagram (EED)

EED As a layered set of instance diagrams, the ecosystem ensemble diagrams presented in this chapter lend themselves to usage and discussions similar to the data flow diagram (DFD)(Yourdon, 1989). Ensembles as a construct for system analysis and design have their roots in (Wirfs-Brock and Johnson, 1990)(de Champeaux, 1991). (Yourdon, 1989) provides a balanced view of the system which is hierarchically consistent with the functional decomposition of the system. Our EED technique changes the focus of decomposition from the functional breakdown, to a perspective which considers regions of interest, data sensitivity, and business transactions and zone vulnerabilities. (Wirfs-Brock and Johnson, 1990)(de Champeaux, 1991) provide an object-oriented view of the system, in its context, with a view of understanding clusters of objects and their roles and responsibilities to the consumers of their behavior. Our approach extends this point of view to multi-stakeholder with computer-proxies who engage in digital transaction, and the risks and rewards that surround their interactions

3) The CPE-Threat Model

(Nazarenko and Safdar, 2019) present a survey of privacy and cyber-security issues in CPS. This survey presents a thorough and systematic overview of the literature. A taxonomy of cyber-attacks is presented, which includes denial of service, eavesdropping, and malware. (Nazarenko and Safdar, 2019) also touch on physical issues of safety and on distribution and deployment concerns, reflecting the physical and geolocated nature of these SoS. The survey suggests that a compositional and encompassing model, such as we propose in the CPE attacker, is still not present in the literature. Our impression is that the domain is biased towards cyber taxonomies. We propose to not limit the effect of the threat to the physical, or the time/sensor-perceptual. Our command and control attacker and the ability to animate attacks using EoS-UML seems to fill a gap in the literature.

3.1.1.2 Exclusions

1) **Methodology.** This thesis does not focus on the methodology to be used to construct the ecosystem. This thesis posits that EoS-UML EED can be incorporated into existing Systems and Software Engineering methodologies including UML processes, Prototyping, Agile, Iterative Development or Waterfall Processes. (Laman and Basili, 2003) provide a well-cited overview of the history and evolution of waterfall, iterative incremental development (IID) and prototyping

methodologies. Process for SoS and CPS follow suit. Example in the literature can be found in (Nadira, et al. 2020)(Bonci, et al. 2018).

2) **GRL.** A necessary future roadmap item in the exploration of EoS UML for prototype evaluation, and operations-time analytics is the incorporation of Goal oriented Requirement Methods. GRL can function along with EoS-UML during requirements elicitation and calibration for the specification of numeric rules and objectives functions (Anda and Amyot, 2019).
3) **Game theoretical and multi-objective optimization.** As mentioned above, possibility to do convergences analysis on deployed collaborative algorithms may be done using game theory. Specifically, collaborative interactions may be analyzed using Stackelberg equilibrium and the competitive (adversarial) interactions between adversary (the CPE-Attacker) and the Challenger (the deployed algorithm and its infrastructure) may be analyzed in terms of Nash equilibrium (Basar and Cruz, 1981). In Chapter 6 trend analysis is demonstrated on a Monte Carlo model.

3.2 Building Blocks

3.2.1 Modeling Notation

This chapter uses UML notation to communicate ecosystem concepts. The UML is a widely used notation to describe system composition and behavior. It emerged in the mid-90s through a collaboration between Grady Booch, Ivar Jacobsen, and James Rumbaugh (Booch et al. 1999), and it defines the notation for modeling systems. Different types of UML diagrams are offered, notably, class, sequence, use case, and

collaboration diagrams. In addition, UML profiles can also be created. In this thesis, UML class diagrams are used to describe the components of the ecosystem and to model credential and point-of-presence relationships.

The main parts of a UML diagram are classes and their relationships. Classes (represented as rectangles) define concepts, and a class may have attributes and operations listed within the rectangle. A system is created by a set of classes and instances thereof, operating together in relationship. On the UML class diagram, a relationship can be binary or n-ary. A binary relationship between classes is drawn as a line between the classes. Binary relationships can have adornments based on the characteristic of the relationship; for example, cardinality can be denoted by numeric adornments, containment by a small diamond adornment, and inheritance (subtyping) by a triangular adornment.

A domain-specific profile may be created using the UML. Indeed, such a profile is discussed for the CPE in this chapter. Profiles are defined using domain-specific stereotypes and mapping them back to a core UML type. This permits a base type to gain semantic tags and semantic and visual adornments according to the profile. Profiles have existed since UML 1.0 and have been revised in UML 2.0 (Selic, 2005); however, while expressive and useful, they are not widely used by industry practitioners. UML profiles have been summarized in the literature, including in (Fuentes-Fernandez and Vallecillo-Moreno, 2004)(Selic, 2007).

3.3 Conceptual Models and Application

This section presents the key concepts in our ecosystem approach to SoS and its application to CPS. The Ecosystem of Systems UML profile (EoS-UML) and the CPE threat model (CPE-TM) are described, and examples are included that relate the two concepts to SoS and CPS concepts and the ATBS domain.

3.3.1 Ecosystem of Systems UML Profile

The ecosystem of systems UML profile (EoS-UML) provides a framework that can be used to model various ecosystems. It is discussed in generic terms with a UML metamodel from which a profile is defined and then specifically applied to the ATBS mTA application.

3.3.1.1 The EoS-UML Metamodel

Figure 5 presents a partial view of the EoS-UML metamodel. The metamodel groups the $structural$ elements and $active$ entities that generically describe an ecosystem. Each entity in the partial view may act as a façade to a cohesive set of concepts within it. The elements in the $structural$ package represent the terrain, environment, infrastructure, and pathways of the ecosystem, while the entities in the $active$ package represent the stakeholders, their devices, certifications, and transactions within the ecosystem. These entities and elements provide a general pattern language for ecosystems, and together, they are useful in defining and exercising applications on the ecosystem. When applied to the entities of a concrete ecosystem, they impart semantics, roles, responsibilities, and an intuitive ensemble visualization.

Entities in the $structural$ package include the $EcosystemOfSystems$ (EoS), $Concourse,$ and the $Region\ of\ Interest$ (or "cell") (RoI). The EoS is traversed by the $Concourse$ which routes processing through aggregative and recursively composed RoI. Moreover, the EoS encapsulates the terrain; it contains an environment and provides infrastructural services.

The EoS' arena of operations, itself, is an RoI, which may be further broken down into subregions. An RoI is a grouping construct that gathers material, human, or service resources; it may occupy a physical area in a mostly ephemeral manner, having a dynamic or short-lived point of presence, which encompasses mostly virtual elements.

The $\boldsymbol{concourse}$ is the set of possible paths between $cells$. Paths may, again, be physical constructs, but they may equally be virtual concepts. The $structural$ elements are activated through the dynamics of the entities in the $active$ package.

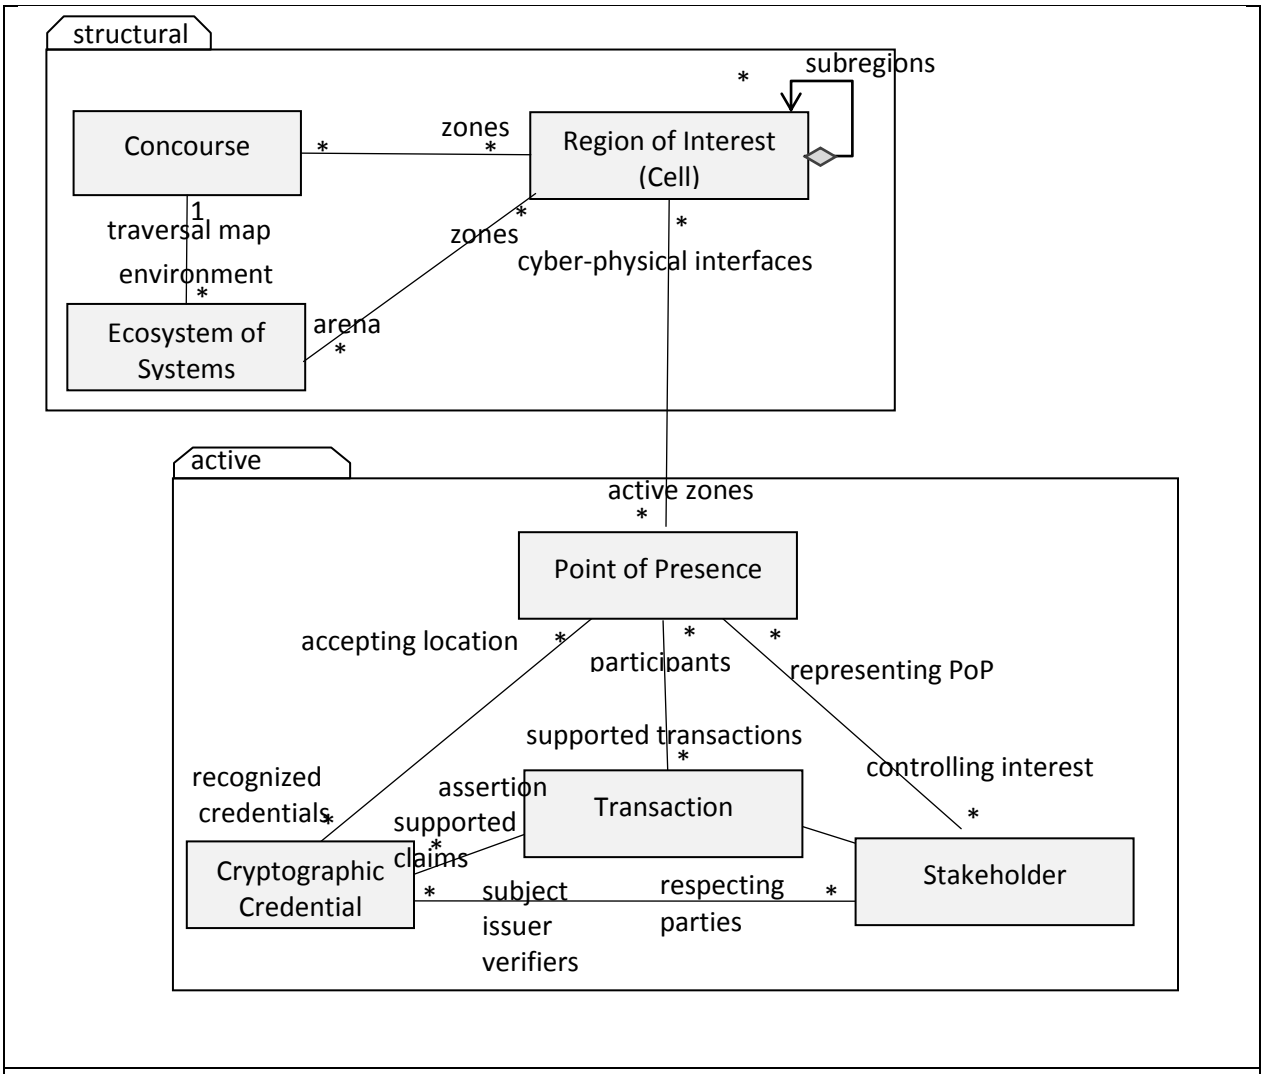

Figure 5) The Ecosystem of Systems UML (EoS-UML) Metamodel

Entities in the $active$ package include the $Point\ of\ Presence$, $Stakeholder$, $CryptographicCredential$, and $Transaction$. Entities in the $active$ package engage, interact with, manage, and (occasionally or habitually) also attack the ecosystem.

A $\boldsymbol{Stakeholder}$ is a meta-concept representing a sentient human our organization of such will well-defined objectives, assets and intent, for example a service provider company, a government agency, or a citizen. A $Stakeholder$ is a goal-driven, risk-averse entity, which conducts business transactions using a point-of-presence device,

and may issue hold and/or verify credentials. In the ATBS mTA setting, the stakeholders are the travelers, the immigration authority, the airline, and the border services authority.

The $\boldsymbol{Point\ of\ Presence}$ is a computing device owned by a $Stakeholder$ and acting on behalf of that $Stakeholder$ in the terrain, conducting transactions, sensing, evaluating, and actuating the subject and environment context. In ATBS, both authorities and subjects have PoP devices. The traveler's PoP is a smartphone configured with a wallet and transaction management application, and the immigration authority's PoP consists of a web interface and a supporting server system, which is rendered on the traveler's smartphone and conducts issuance functions for the mTA. In this thesis, the airline and the border services agency are verifiers of the mTA; their points of presence are kiosks, cameras, microphones, turnstiles or gates, signage, and employee-facing handhelds and desktop systems.

Through their points of presence, stakeholders engage in $\boldsymbol{Transactions}$. Transactions are business-level functions that subjects and authorities engage in for the request and granting of service. In our ATBS setting, some transactions include "Traveler application for mTA," "Pre-board screening request," and "Request to enter country of destination."

The $\boldsymbol{CryptographicCredential}$ is a bundle of certified attributes about a subject; it is issued by a trusted stakeholder in an <<*Issuer*>> *role*, managed by the *subject*, and redeemed by to respecting *Verifiers*, in support of some value-interchange transaction. These three roles (*Issuer, Subject,* and *Verifier)* map to the standard credential roles as

in (Brands, 2000)(Camenisch and Lysyanskaya, 2001). Stakeholders engaging in a transaction assume one of those roles for the transaction. A transactor may have honest or nefarious intention during a transaction.

3.3.1.2 The UML Profile

The metamodel presented above can be used to define a UML profile, which can be applied to ecosystems of various domains.

The EoS-UML profile is defined by creating a stereotype for the concepts in the EoS -UML metamodel, extending the UML type to which it can be applied, and optionally adding tagged values and graphical icons. A partial profile for the EoS-UML is illustrated in Figure 6.

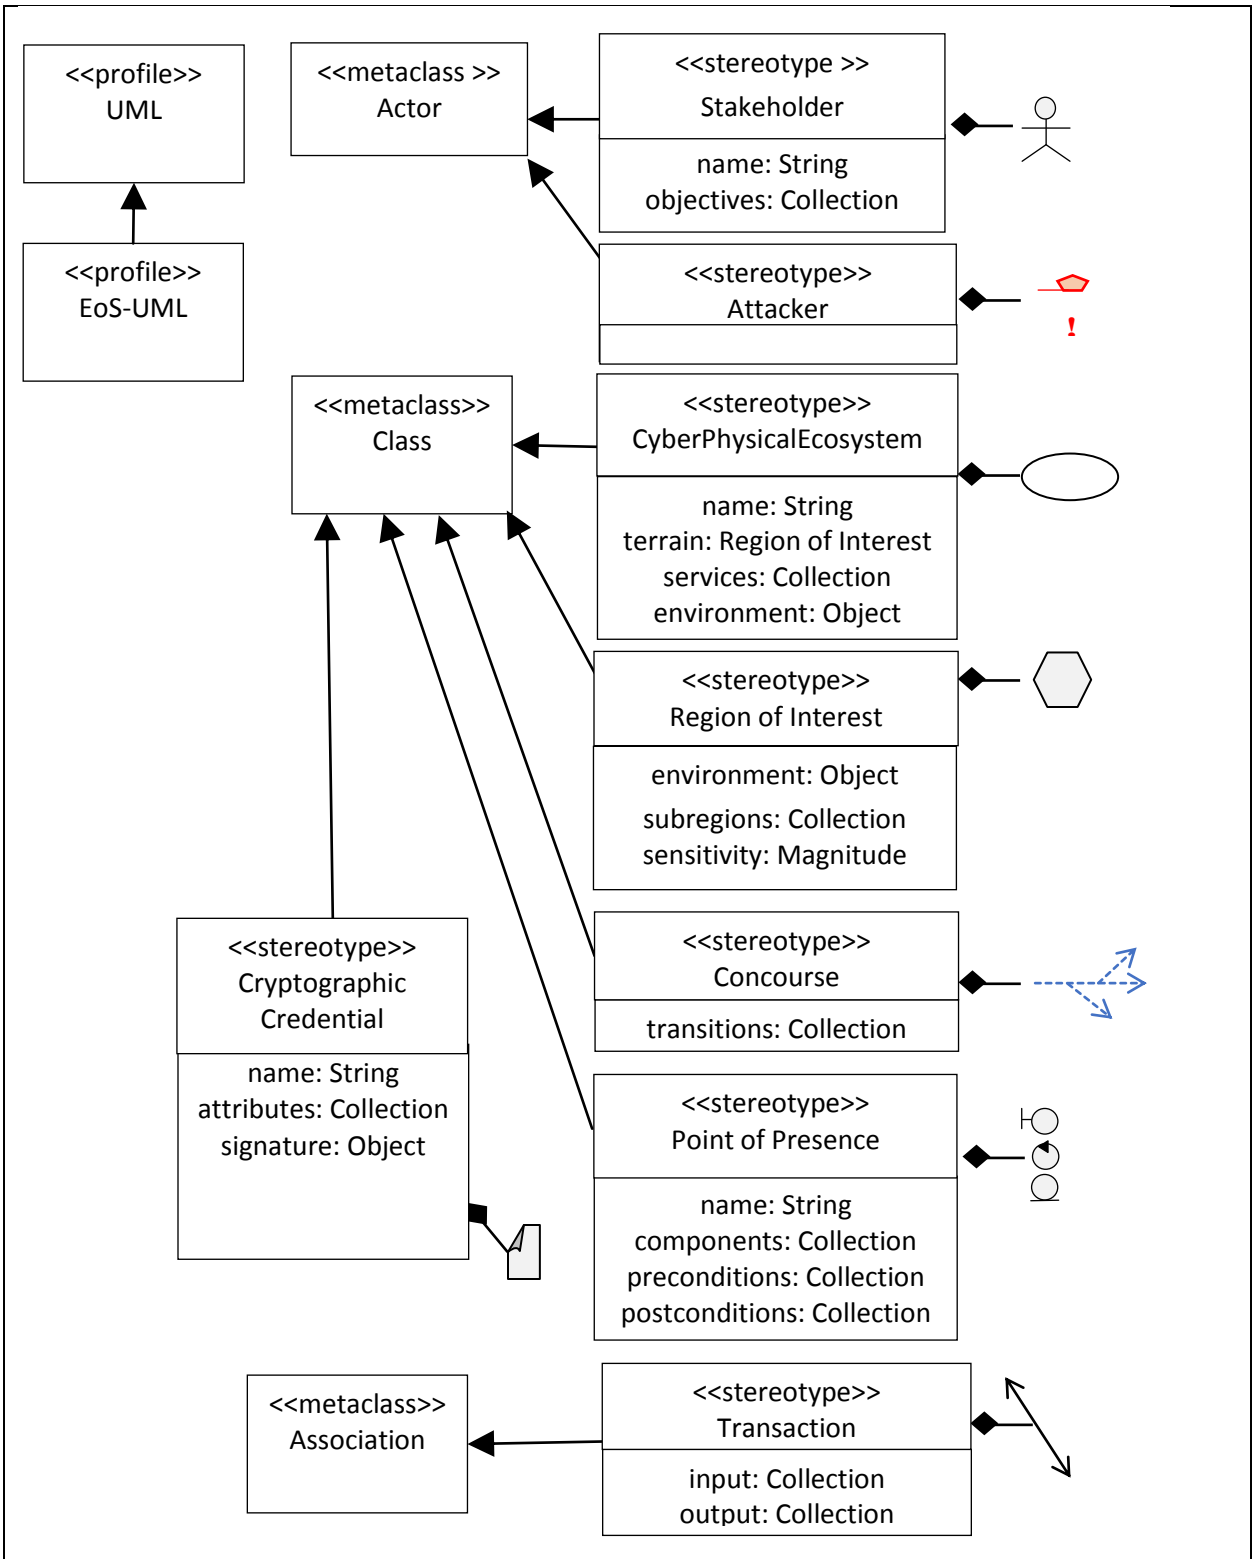

Figure 6) The Ecosystem of Systems Unified Modeling Language Profile (EoS-UML)

Figure 6 defines eight stereotypes. $Stakeholder$ and $Attacker$ extend the $Actor$ metaclass; $Ecosystem$, RoI, $Concourse$, and $Point\ of\ Presence$ extend $Class$; and $Transaction$ extends $Association$ from the base UML system.

A $\boldsymbol{Stakeholder}$ has a name and a set of goals and uses the default actor icon. An $Attacker$ can be applied to a $Stakeholder$ and a $Point\ of\ Presence$ (recall the EOS-UML metamodel in; Figure 5). This will be further discussed §3.3.5 on the CPE-TM (see p. 86). A specialized adornment has been provided for an $Attacker$– either a red hat or an exclamation mark can be added to corrupted resources in diagrams.

The $\boldsymbol{Ecosystem}$ stereotype has tagged values for the ecosystem name, the ecosystem terrain, and the ecosystem services. The terrain is an RoI. In ATBS, services include transcript, directory, and clock services. The ecosystem is represented as an oval in diagrams, with tagged values specified in an associated note.

The RoI stereotype has tagged values for its environment, its cyber or physical damage sensitivity, and its subregions. The environment captures engineering-relevant attributes such as threat alert level and congestion in the case of an airport, for example. Cyber or physical damage sensitivity is a magnitude estimate of the damage per resource of the protected elements in the zone. This rolled-up magnitude assists authorities in determining their thresholds and risk appetites for error rates in the zone. The subregions are a collection of (possibly nested) $RoIs$ that constitute the instance. A custom RoI, for example, may contain three subregions – one each for primary inspection, secondary inspection, and payment transactions.

The *Concourse* instances connect RoIs. A concourse is a lifecycle or state machine through the regions. It can represent a number of things, including the collection of all possible routes or a lifecycle of ecosystem transactions; however, it may be easily and correctly considered as an application. It has a tagged value of transitions and a graphical icon denoting a network of paths.

The *Point of Presence* is a compute-technology object that is assigned a name, components, and contractual collections for pre- and post-conditions and invariants. The components may include a display, sensors, actuators, an application server, and a database. This object is given the diagrammatic icon as in the boundary-control model (BCM) pattern (Jacobson, 1992). These icons may be distributed within the region, in the cloud, or alongside a stakeholder to communicate the nature of their architectural placement. The pre- and post-conditions and invariants are contracts guaranteed by the PoP to the concourse and stakeholders. These conditions together form a transaction contract, which will be discussed further in Chapter 7 “Ecosystem Design by Smart Contracts (DbSC)”. Example pre-conditions include “Subject must have been authenticed within t certainty” or “Subject must be >= eighteen,” and example post-conditions include “Identity-fraud risk residue will be below 5% at a 95% confidence interval” or “identity assertion is be vetted at an information assurance level of 3.” Invariants are predicates applicable on a contracted procedure that are true throughout its execution, for example “Passport photos are not saved to backing store”.

Cryptographic Credential represents issuer-certified data that may be held and redeemed by verifying organizations in value transactions. The cryptographic credential

has a name, attributes, a signature, and a graphical icon corresponding to a rectangle with a folded corner.

A *transaction* encapsulates the interchange which occurs between transactors of value for consideration. All transactions comprise an electronic component, some transactions occur in person with one or more stakeholder representative present in-person on the terrain, at the point of presence. Often, however, one or both stakeholders engage in the transaction in a position remote physical terrain, present through digital means. Transactions have both cyber and physical qualities.

This partial discussion informs the threat model, the issuance and verification workflows of Chapters 4 and 5, and the collaborative computing architecture of Chapter 6. Many additional stereotypes and attributes are possible; these can be specified on a domain basis and are left for future work.

3.3.2 Ecosystem Ensemble Diagrams: Instantiating the EoS-UML on ATBS

This section applies the EoS-UML to the ATBS mobile travel credential scenario. First, the ecosystem ensemble diagram is presented to provide a high-level context of the system. Then, a Level-1 concourse diagram is presented. In addition to demonstrating the EoS-UML.

3.3.2.1 Extending the UML Profile

As discussed in the previous section, the EoS-UML is generic. It can be extended in a domain-specific manner with instance icons representing the particulars of the domain. As an example, consider Figure 7 which presents an EoS-UML extension for ATBS,

adding some custom stereotypes with specialized attributes and graphical representations.

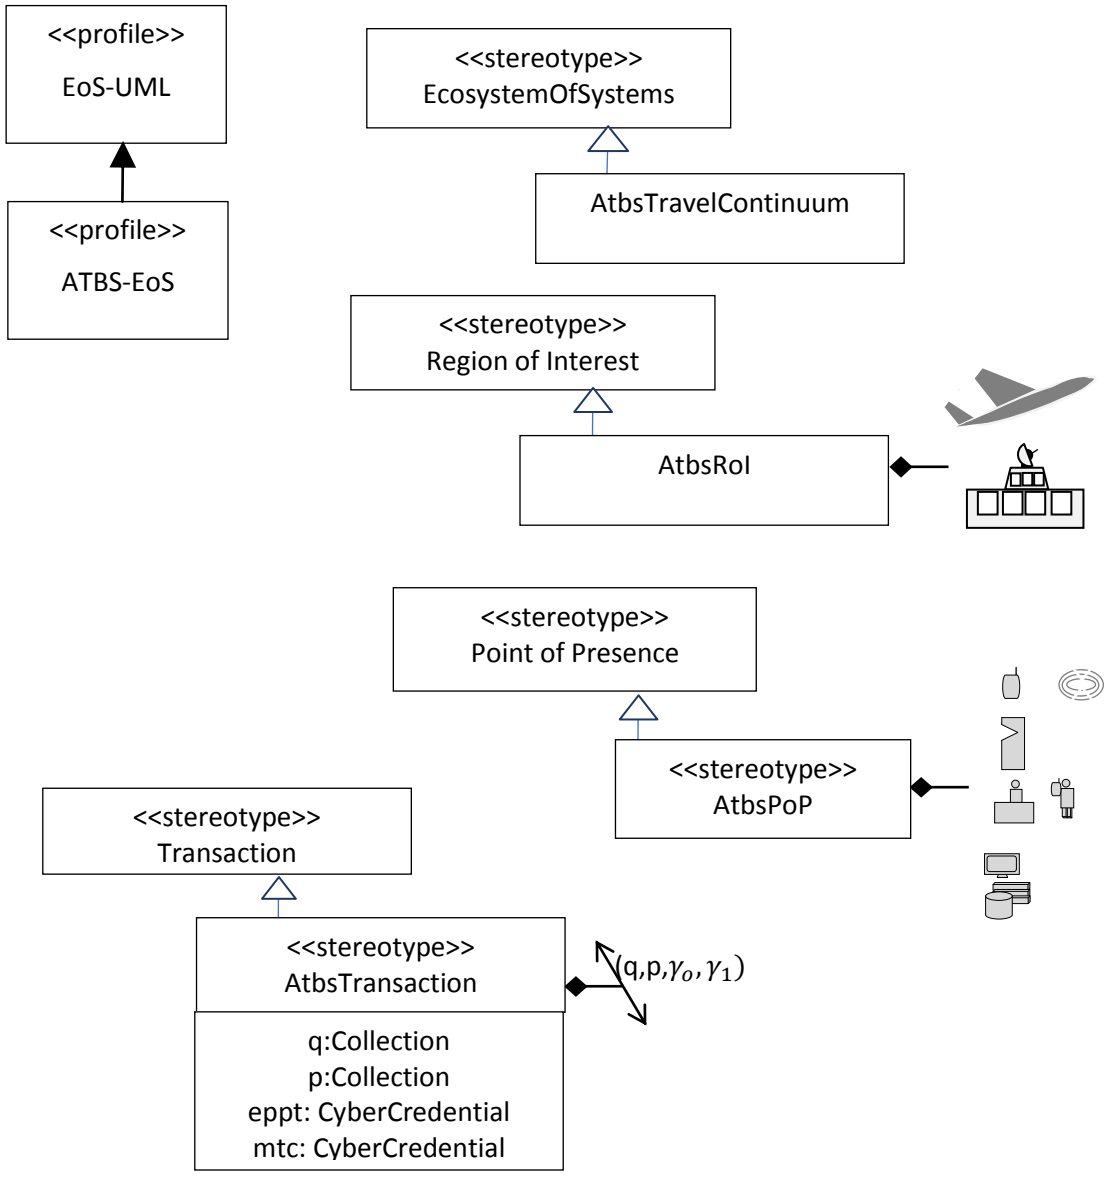

Figure 7) Extension of EoS-UML with ATBS Concepts

With the exception of «Actor», in UML 2.0, stereotypes cannot extend other stereotypes. They can, however, “specialize” or “generalize” – a relationship shown with the inheritance association, as above. Genericized Magnitudes and Collections are assumed and expressed in the interface. The concrete specification occurs when the EED is instantiated on a zone of interaction (the ecosystem itself, or a sub-RoI) and specific transactions, relationships and attacks are elaborated refining the collection, transaction, attribute, magnitude, lexicon and over-the-air requirements.

The profile in Figure 7 illustrates generalizations of a $Transaction$, $Point\ of\ Presence$, and $Region\ of\ Interest$. ATBS could add a kiosk and an attended desk PoP for the airline and the border services, a smartphone for the traveler, and server equipment to communicate with corporate stakeholders. A general mTA transaction can be defined, adding a request, a response, and e-passport and mTA input credentials. Specific extensions could be defined for «$Airport$» and «$Airplane$» RoIs.

3.3.2.2 Level 0: The Ecosystem Context Diagram

Figure 8 presents a context-level ZED in which the subject and service providers are shown at the periphery of the CPE, interacting in credential transactions.

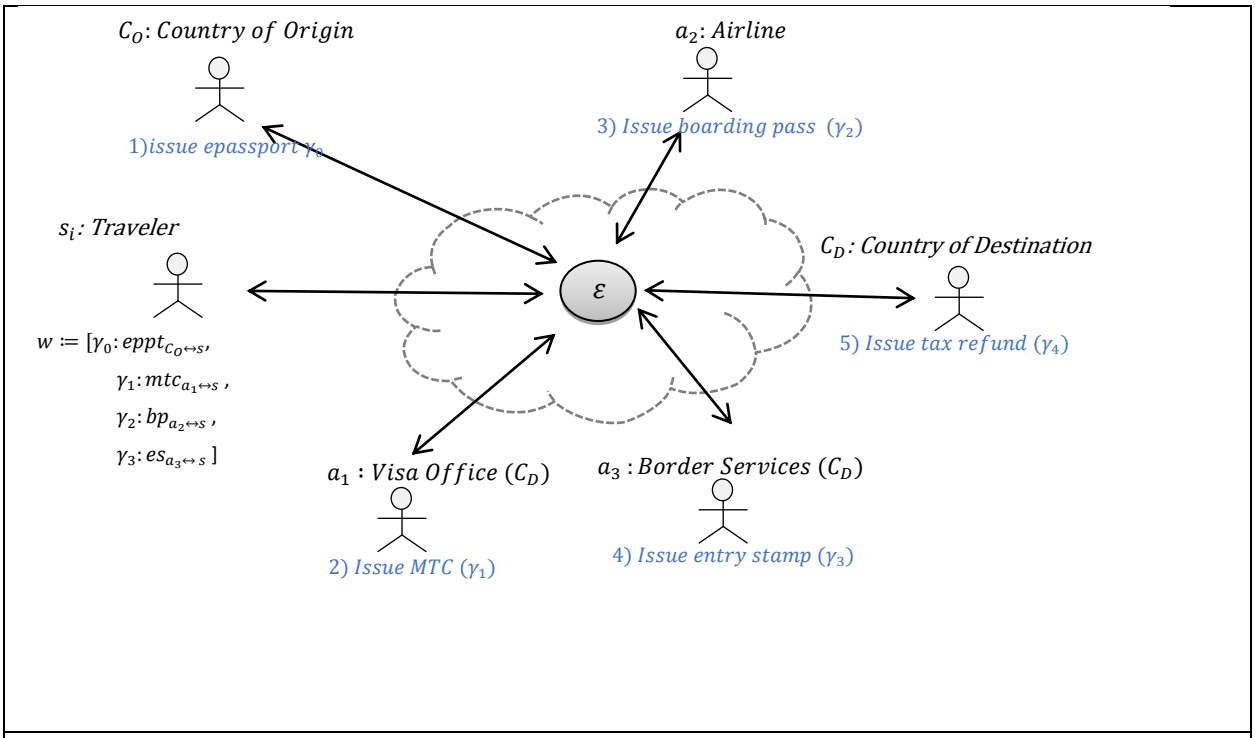

Figure 8) ZED Level 0 – Ecosystem Ensemble Diagram for ATBS

In the context-level ZED of Figure 8, the subject, a traveler, and five authorities are depicted on the periphery of the ecosystem. Arrows between these stakeholders and the ecosystem indicate transactions and suggest a lifecycle. The names of the transactions are also numbered to further communicate the lifecycle. In Figure 8, the subject participates in each numbered transaction, accumulating credentials in wallet w over the course of a series of transactions with the various system authorities. Wallet w is part of the PoP of s_i. The wallet is optionally shown to help to communicate the transaction lifecycle.

The transactions on the ZED tell a story. In consultative situations, in which the goal is to elicit requirements from clients, this ability to "tell a story" can help to drive discussion which uncover important business rules and domain subtleties. In Figure 8,

the subject first obtains an ePassport γ_0 from C_0, the citizenship authority of their home country. In preparing for travel, the subject then obtains an mTA γ_1 from a_1, which is the visa authority of the country of destination C_{D}. At the airport, the traveler obtains a boarding pass γ_2 from the airline a_2 after successful verification of their travel documents γ_0 and γ_1. Then, in the destination country, the traveler obtains an entry stamp γ_4 after successful verification of their travel documents γ_0 and γ_1 and entry declaration. At the end of their trip, they may submit tax receipts and obtain a tax refund γ_5.

This thesis trims the scope of Figure 8; we focus only on the issuance and verification of γ_1 in conjunction with the travel of s_{i}, who is assumed to hold γ_0. The thesis does not examine the issuance of other credentials. The assumption is that the modeling and design of those credentials will follow similar techniques.

3.3.3 Level 1: Contexts and Main Concourse

Figure 9 presents the next level of detail, within a scope limited to the issuance and verification of γ_1 by authorities a_1, a_2, a_3.

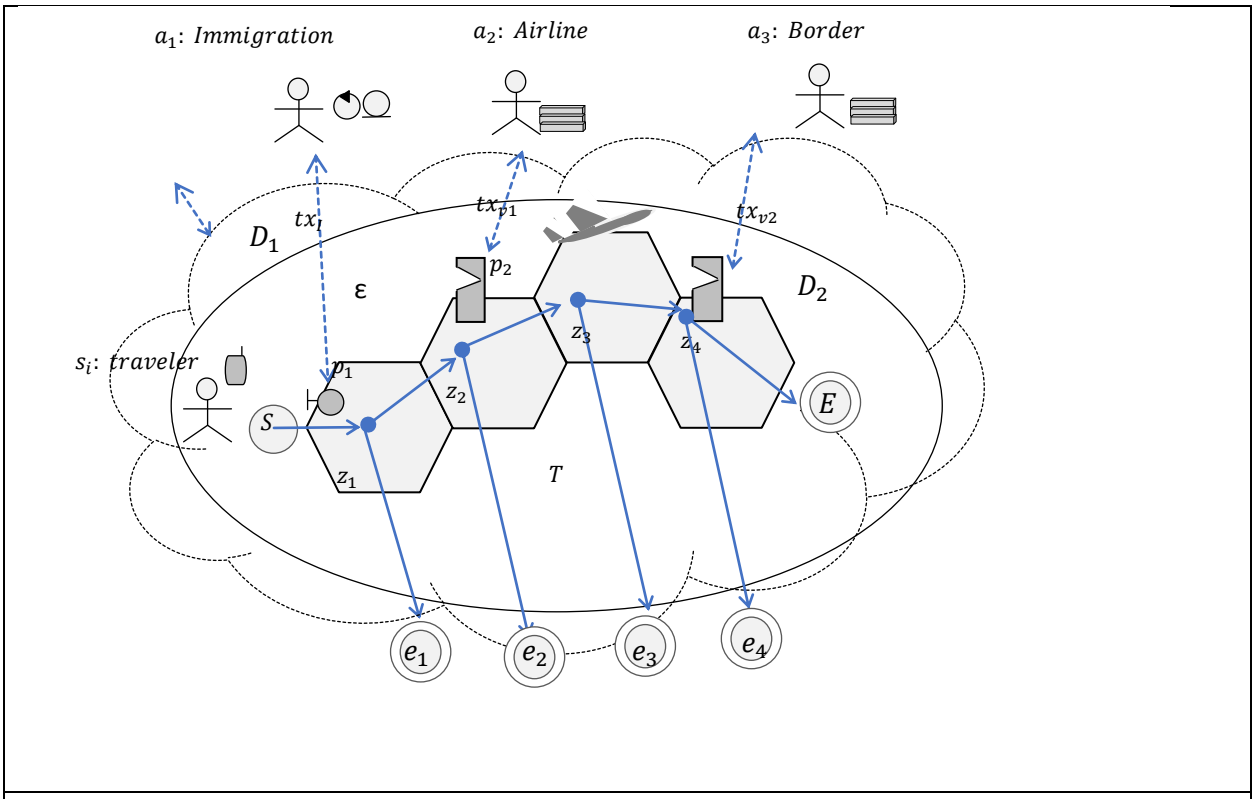

Figure 9) EED Level 1– The Concourse and Transactions

Figure 9 presents the Level-1 ecosystem ensemble diagram (EED), which depicts the ecosystem's stakeholders, the contexts they control, and the connected concourse graph. The following aspects in particular must be noted about Figure 9:

1) **One Ecosystem with Three Services.** Figure 9 contains $Ecosystem$ ε with three instances of *Service*: A *Transcript* and two *Directory* objects(T, D_1, D_2).
2) **Three ROIs.** $Ecosystem$ ε contains four instances of $RoI(z_1, z_2, z_3, z_4)$. These RoIs are controlled by three authorities (a_1, a_2, a_3).

One Concourse. $Ecosystem$ ε contains one concourse with a specified start and end node, three checkpoint nodes, and three early exits. The start and end nodes represent the normal flow through the concourse, and the early exits are defined processing paths occurring in exceptional situations. (The

concourse may be interpreted as an application or an object lifecycle; these terms may be suitable in the long run, but for now, we use “concourse”.) The UML profile does not suggest representation for a concourse: a state machine approach is used in Chapter 6, however a decentralized orchestration platform would likely be used in operations.

3) **Three Transactions**. Three transactions are shown – one between each authority and the ecosystem. These are labeled tx_I, tx_{v1}, and tx_{v2}, representing the mTA application, these may represent the request for an mTA, the request for a boarding pass, and the request for admission, respectively. Importantly, this demonstrates how business-level transactions encapsulate cryptographic credential operations. Here we have an issuance and two verification operations. As the project team refines their understanding of the qualities of honest and nefarious value interchanges in a region, additional candidate transactions and credential requirements may be brought to light and included or omitted from scope as appropriate for ecosystem objectives and required scaffolding for its long-term mandate.
4) **Four Points of Presence**. Each of the four stakeholders are given a $Point\ of\ Presence$. Different icon sets are displayed to represent in-situ elements versus computing resources that are housed externally to the CPE. The compute stack for a_1 shows the familiar OOSE icons (Jacobson, 1992). The diagram assumes that the UML profile has been extended to include icons for smartphones, kiosks, and the server stack (see §3.3.2.1 and Figure 7).

3.3.4 Level 2: Subregions and a Full Concourse

Figure 10 presents a Level-2 decomposition of the ecosystem into its subregions and illustrates a more detailed concourse than the Level-1 diagram in Figure 9 with stakeholder transactions. The labels on the elements are symbolic for clarity. Subjects $\{s_1, \ldots s_n\}$ enter at the start node S with a main flow crossing contexts z_1, z_2, z_3, operated by authorities a_1, a_2, a_3 towards their stated objective, namely, exit node E. Alternative flows can occur at checkpoints p_1, p_2, and p_4, which can result in stable alternative exits identified at e_1, e_2, and e_3.

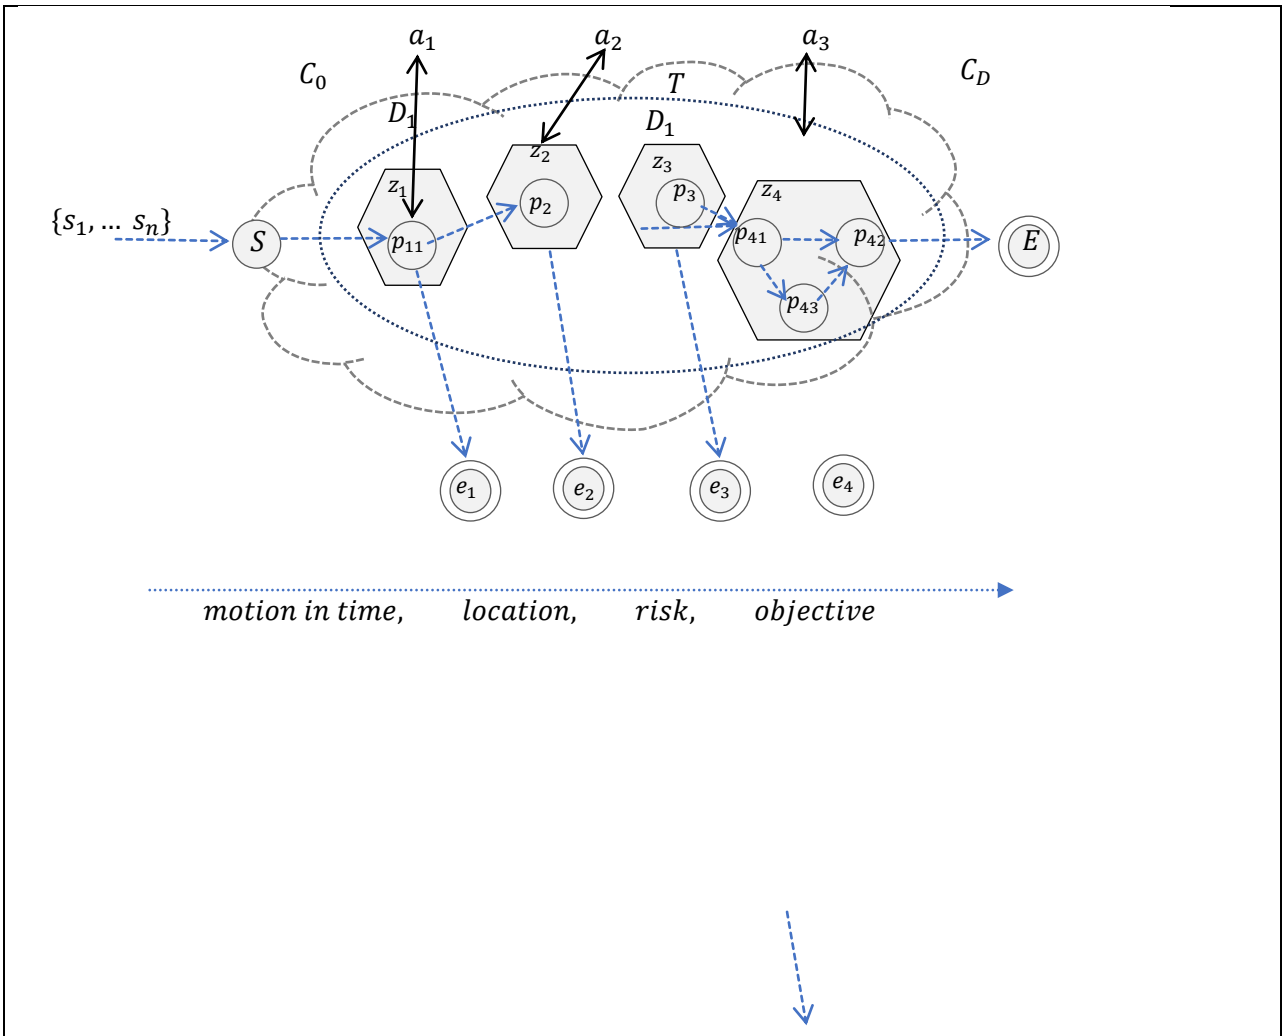

Figure 10) EED Level 2 – Recursive Decomposition with Balancing

The concourse presented in Figure 10 adds details the mTA workflow presented in higher level ZED. Notably, the border security zone z_4 is contains three checkpoints p_{41}, $p_{42,}$ and p_{43}, implying subregions, any of which may also be recursively composed. The checkpoints contained in z_4 could be given business functionality, such as "automated processing region," "manual exception processing," and "final triage." The automated processing region might be made up of a web-delivery questionnaire, a collection of

kiosks, a bank of e-gates, and an officer hand-held system that receives alerts for required assistance. A checkpoint typically implies a processing region. The subregion protected by the contained checkpoint may be shown as a hexagon with adornments if helpful to the audience at hand.

Compositional Refinement and Zone Ensemble Diagrams. During the compositional analysis and refinement of the transactional value interchange which occurs in an ecosystem and its zones, the ecosystem designer may choose to focus on the zone rather than attempt to capture and rationalize the detail to the high-level ecosystem context. In this case, the focused diagrams may be referred to as Zone Ensemble Diagrams (ZED) rather than Ecosystem Ensemble Diagrams. It is not incorrect to refer to the top-level context diagram as a ZED also. This is due to the fact that the ecosystem's arena of operation (its "terrain") is define to be a region of interest in EoS-UML.

Progression Through the Concourse. The left-to-right motion along the arrow shown in Figure 10 can be used to informally communicate a number of things: for example, the passage of time, the change in location or context, the progressive motion of a subject toward an objective, increased confidence in the results of the risk screening process, increased penetration into the ecosystem, accumulated cost, or accumulated error. This left-right positioning can be used to increase the story-telling qualities of the EED.

The progress of a particular subject through the concourse reflects the progression of the system's work in having routed, categorized, and processed that subject. At checkpoints, electronic interviews are conducted using subject and authority PoP

equipment to exchange data which informs decisions about the next step in the process for each subject. Decisions are sensitive to location or context as well as environmental and subject-specific attributes. This is in fact a distributed classification performed on a continuum with a distributed configuration that includes local and system-wide costs and/or benefits.

From the perspectives of an authority, a subject, and the ecosystem itself, the sequential “progress” of a subject through the continuum may be evidenced by accumulated transactions, which take the form of cryptographic credential receipts, routing slips, and transaction mementos in corporate datastores, in subject wallets, or on the transcript.

3.3.5 The CPE Threat Model

An ecosystem perspective also pays off when applied to threat modeling. A traversal through the cyber-physical space involves a series of actions by stakeholders and their systems in time and physical space in the ecosystem. These aspects must be reflected in the threat model. The proposed CPE-TM generalize CPS threat models and the standard threat model used in cryptography in that the attacker $\widetilde{\mathrm{A}}$ incorporates human, machine, location, and temporal elements.

3.3.5.1 The Impacts of a CPE Attack on Various Stakeholders

Loukas’ definition of a cyber-physical attack includes the criterion of the attack having a physical effect (Loukas, 2015). However, categorization based on observable physical effect may be problematic. A CPE has numerous complexities and vantage points. These include multiple stakeholders, continuous time, and an evolving contour or

contrasting objectives to name a few. The symptoms of an attack may not be evident. The observable symptom may be an unwanted by-product of an attack. The adversary may seek stealth and dissimulation. Observability may be a symptom of a failed attack. Similarly, defining a cyber-physical attack as necessarily having a *physical* effect may be limiting. The key to detection of a cyber-physical attack and successful remediation may lie in intercepting the attack while it is still in its non-physical manifestation. Finally, we propose that the criteria for a CPE attack should consider not only the effect but also the stimulus. A cyber-physical attack has a complex attack surface with origin points across the ecosystem terrain. Classification of attacks should thus benefit from identification of the origin and nodal points of the attack vector and its trajectory.
We propose that the criteria for a cyber-physical attack taxonomy should include stakeholder objectives, the time- and location-specific region invariants, and the attack origin and propagation path. An attack is implemented by the attacker using possible combinations of cyber impulses and physical impulses in a concerted, chronological, and geolocated manner toward possibly hidden attacker goals. Similarly, its effects are manifested in a combination of cyber and physical symptoms, which may (or may not) be detectable.

The symptoms of an attack will be perceived differently by different stakeholders. The manifestation of these symptoms may also differ depending on the cyber-physical reach of the various stakeholders in the ecosystem.

The impact of an attack may be cyber-based or physical and may be felt differently by various stakeholders in the ecosystem. It is important to distinguish the attack's impacts from the attacker's specific goals. An attacker may achieve their goal without

apparent symptoms felt by the stakeholders. This is not to say that the attack does not have an impact on the stakeholder(s), only that it has not (yet) been felt. Types of attacks include credential lending or credential theft combined with biometric substitution and attribute falsification.

3.3.5.2 Attacker Description

The adversary $\widetilde{\mathrm{A}}$ is made up of three components: $\widetilde{\mathrm{Ac}}$, the command and control attacker; $\widetilde{\mathrm{Ap}}$, the attack platform; and $\widetilde{\mathrm{At}}$, the terrain attacker, thus $\widetilde{\mathrm{A}} = (\widetilde{\mathrm{Ac}}, \widetilde{\mathrm{Ap}}, \widetilde{\mathrm{At}})$.

The first component of $\widetilde{\mathrm{A}}$ is $\widetilde{\mathrm{Ac}}$, a malicious human entity that defines the attack objective and leads its execution. The $\widetilde{\mathrm{Ac}}$ component uses specialized intelligence and judgement to dynamically resolve choices and sequence actions as necessary. The second component of $\widetilde{\mathrm{A}}$ is $\widetilde{\mathrm{Ap}}$, the algorithm platform library available to the attack. The third component of $\widetilde{\mathrm{A}}$ is $\widetilde{\mathrm{At}}$, which consists of a troop of malicious human actors ("*field operatives*") and/or their computer platforms located throughout the ecosystem. At the start of an attack, $\widetilde{\mathrm{Ac}}$, $\widetilde{\mathrm{Ap}}$, and $\widetilde{\mathrm{At}}$ are instantiated with objectives, algorithms, and personnel. The size of $\widetilde{\mathrm{At}}$ changes over time; $\widetilde{\mathrm{At}}$ grows as honest actors in the field (or their computer equipment) are corrupted.

Furthermore, $\widetilde{\mathrm{Ac}}$ is a goal-driven human attacker who observes and acts on the ecosystem using a distributed workforce of personnel and resources ($\widetilde{\mathrm{Ap}}$ and $\widetilde{\mathrm{At}}$). This type of attacker grows in strength by corrupting additional human and computer resources over time, adding to $\widetilde{\mathrm{Ap}}$ and $\widetilde{\mathrm{At}}$. Through expansion, $\widetilde{\mathrm{A}}$ may gain either deeper penetration into ecosystem contexts, additional protected data, algorithm or a richer overall vantage point on the ecosystem attack surface. The growth of $\widetilde{\mathrm{Ap}}$ and

$\widetilde{\mathrm{At}}$ can occur in stages, proceeding, for example, via a social engineering attack, the infection of computation resources, or the diversion of human attention and objectives, and spreading to the relationships and sphere of interest of the human subject. As the actor becomes corrupted, they are inducted into $\widetilde{\mathrm{At}}$ (as an “agent”). Agents provide distributed reconnaissance intelligence, and attack capabilities. Using a military analogy, $\widetilde{\mathrm{Ac}}$ performs the command-and-control function (C2), while $\widetilde{\mathrm{Ap}}$ are the resources, and $\widetilde{\mathrm{At}}$ are the troops on the ground.

The attacker’s goals may include information theft, context infiltration, and sensor or actuator hijacking. In general, an attack aims to compromise smooth ecosystem functions by breaking security, privacy, and operational properties.

The physical fleet $\widetilde{\mathrm{At}}$ provides the attack model with human capability and position in time and space. The algorithmic base $\widetilde{\mathrm{Ap}}$ is polynomially bound and may be used in active manners to attempt to defeat cryptographic security, create spoof biometric images, establish surveillance channels, or run other algorithms required by the attack.

Troops in $\widetilde{\mathrm{At}}$ consist of a human element and possibly a computer component. This is parallel to the human-to-machine relationship evidenced in the ecosystem in which stakeholders in the ecosystem act within the ecosystem through their connected computing devices. This is true of the adversary as well: when attacker $\widetilde{\mathrm{A}}$ corrupts an honest participant, the computer, the actor, or both may be corrupted. Corrupting a computing device places it in the control of the attacker, who is capable of sensing the area around the node and actuating according to its capabilities. In many ways, corrupting the actor is a stronger attack, since a corrupted actor can affect the context

with human capabilities (situational awareness, intelligence, and interaction), combined with the computational strength of the device.

At the start of an attack, $\widetilde{\mathrm{A}}$ is instantiated with baseline troops (containing zero or more individuals and machines). Over time, $\widetilde{\mathrm{At}}$ changes as the attack unfolds because $\widetilde{\mathrm{A}}$ corrupts honest participants or machines throughout the CPE. Moreover, as honest participants are corrupted, $\widetilde{\mathrm{A}}$ gains both cyber and physical strength. The capability of a corrupted actor is measured in terms of the actor's presence in-situ, their authorizations, physical potential, and computing resource.

3.3.5.3 Cyber and Physical Components of Platform and Terrain Attackers

Table 5 lists the cyber and physical components that both a human and a computer are assigned in our threat model.

Table 5) Duality of Ecosystem Components

Entity / Component	Human Operator	Machine
Cyber	Electronic data or profiles App personalization Permissions or authorizations Credentials or helper data Biometric images and templates Thresholds	Software stack, communications parameters, open ports, API O/S, Patches
Physical	Individual, representative Goals Vantage point Position	Environment Sensor fidelity Light source reliability

The human and system components of the CPE and the attack model both involve cyber and physical components. Therefore, the threat model must be considered with human and computer components.

In addition, we add the key factors of time and location to these components. An attacker's location may allow the compromise of the environment or system within a protected context. Penetration into a hall containing verification kiosks, for example, opens potential for the Attacker to disrupt biometric operations by saturating (or depriving) the scene with light. The timing of an attack must be considered. Once a context is penetrated and biometric operations are disrupted, a momentary diversion using social media or intercom announcements can further compromised stability of standard operating procedures (SOPs) to enhance the probability of a successful attack.

3.3.5.4 Specifying an Attack: Magic Passport.

Consider the magic passport attack (Ferrara et al., 2014). In this attack a *blend(…)* algorithm is used to combine valid facial images of two or more malicious individuals into a single blended image $\tilde{b}_e$ which is submitted to an issuer to obtain a valid, but defective passport which authenticates successfully when used by either subject at non compromised verification station.

To specify this attack, let the initial state tuple be As_0 = ($\widetilde{\mathrm{Ac}}$, $\widetilde{\mathrm{At}}$, $\widetilde{\mathrm{Ap}}$), where $\widetilde{\mathrm{Ac}}$ is the single controlling attacker, $\widetilde{\mathrm{At}}$ contains the individuals who will be using the passport to authenticate in the feild $\widetilde{\mathrm{At}} = (s_1, s_2)$, and

$\widetilde{\mathrm{Ap}} = \{ issue(s),\ verify(\ \ ),\ M_I(\ \ ), M_V(\ \ ), blend(\ \ ) \}$ contains the algorithms available to the attacker.

The attack proceeds as follows. A blended facial image is constructed as $\tilde{\mathrm{b}}$ = blend($\{\mathrm{b}_1, \mathrm{b}_2\}$) where b_1 and b_2 are legitimate biometrics of s_1 and s_2 respectively. Blended image $\tilde{b}$ sent to the issuer along with required supporting attribute X to obtain a legitimate passport $\widetilde{\mathrm{ppt}_{12}} = a_1.issue(\tilde{b}, \mathrm{X})$. Malicious individual s_2 produces the compromised document for verification at checkpoint $\mathrm{p_i}$. This verification occurs in $\mathrm{p_i}.verify(\widetilde{\mathrm{ppt}_{12}}, s_2)$ in which the kiosk reads $\tilde{b}\ from\ \widetilde{\mathrm{ppt}}$ and verifies it against fresh biometric b_v using biometric matcher $M_V(\tilde{b}, b_v)$. The performance of $\tilde{b}$ against a valid biometric from s_2 in M_V is known by the attacker. Thus the checkpoint verification will succeed. If the checkpoint determines risk, it can forward the individual processing escalating to a human officer. The human officer may or may not detect that the image on $\widetilde{\mathrm{ppt}}$ is suspect. Thus the Attacker has reasonable chances at infiltrating the protected zone depending on whether escalation to a human-operative occurs.

3.3.1 The Arms-Length Trust Model

The Arms-Length Trust Model underlies the dynamics of the ecosystem's value interchange. The trust model states:

1) **Stakeholders act on their own behalf, towards their long-term objectives**

   Stakeholders are rational and objective driven. They weigh environment and transaction data, choose to engage or not, in the transaction(s) at hand. Stakeholders act on their own interests. Allegiances can form. They may be long- or short-lived. A Stakeholder's role in can chance from "honest" to "malicious".

2) **Rational expectation of satisfaction/consequences of long term participation in the ecosystem**

   Due to the non-determinism that is stitched throughout the ecosystem, it is impossible to predict precise end state. Nonetheless, statistically, given available environment and transaction information, transactors feel rationally justified in their long-term stakes and their probabilities of occurring.

3) **Allegiances may change dynamically**

   Honest and malicious behavior is emergent – swarm-like. For any sentient in the human element of the system, that sentient's alleigiances can change from one transaction to the next. Such change may or may not affect public information in the ecosystem.

4) **Each transaction has inherent risk. and potential damages and/or rewards.**

   The transaction is the vehicle through which value is exchanged and thus helps tokenizes a workday's objectives. Due to possibility of information falsification,

subject or peer misrepresentative, each transaction bears risk. Without engaging in transactions, however, no value is generated. Arms-length participants engage in each transaction selectively -- according to its own (perceived) risk and merits.

5) **The environment and transaction data to assess risk**

Transaction information includes data supplied by the subject, by upstream credential providers, or from the environment. Environment data is contextual, and service oriented, and of a recursive structure, matching in a way the recursive decomposition of resource-sensitivity zones. Environment data may include the purely environmental (such as zone illumination, temperature, and occupancy data) as well as service statistics (vendor service reputation, confusion matrix, and confidence interval expectations, lexicon grammar information values, and service costs) and socio-centric data (peer confidence or reputation assessments, cleansed historic performance data, outlier summaries, etc.).

6) **Not all supplied information is of the same quality**

Transactors and environmental; service may supply inaccurate or false information to a transactor. The onus is on the transactors to assess correctness of supplied information, and accept the decisional risk due to incorrect and uncertain data. In the architecture underlying this thesis, services directory, and the lexicons of information gain, and predictive value provide benchmarks and measures for this uncertainty assessment. In practice, identity proofers and verifiers are often granted an information level based on the diligence they pursue and the level of uncertainty in their attribute attestations.

3.4 Discussion

3.4.1 Elaborating the Metamodel

The metamodel in Figure 5 does not decompose the concepts to their atomic level. Each entity may be thought of as a *Façade* (Gamma et al., 1993)**.**

The *Cryptographic credential*, for example, includes the notion of a *Credential Role,* which in turn covers the concepts of *Issuer*, *Holder,* and *Verifier* and the fact that these are not mutually exclusive adornments. These concepts may be represented as a UML enumeration; however, a stereotype with tagged values may be most valuable and correct. A stakeholder in a transaction offers a business service, but may act in one or more credential roles. For instance, if the primary kiosk offers a "submit declaration" transaction, which requires three credentials, a declaration, an e-passport and an mTA as input, and issues an entry-stamp credential on successful clearance, then this transaction can be annotated as both $Verifier$ and $Issuer$. This treatment is implied in the UML profile with the "in: Collection" and "out: Collection" tags. For an incoming *Cryptographic credential* , the service provider acts as the $Verifier$, and for a return credential, the service provider acts as the $Issuer$.

While associations between actors are not permitted in UML 2.0 (Selic, 2005), subtyping is permitted. Here, the *Stakeholder* entity acts as a supertype. There are different types of *Stakeholders,* including companies, government agencies, their employees, and individual citizens. The specific subtypes depend on the domain and are left unspecified here. Specialized stakeholders can be modeled in different ways in the

metamodel and UML profile. At its simplest, the stereotype $\ll Stakeholder \gg$ may be applied on appropriate entities in the instance model.

3.4.2 Ecosystem Engineering Using CPE-UML Artefacts

A few aspects should be noted about the set of analysis diagrams presented in this chapter. The ZEDs are a layered set of instance diagrams delivering unique value to the software engineering process:

1) **Customizable Look.** The visual manifestation of the ZED can be altered depending on what is most appropriate for the audience. BCM icons may be useful for software engineers, and form factor icons may be more suitable for business stakeholders.
2) **Instance Diagrams.** The use of specialized graphical icons may make these diagrams appear as simple architecture diagrams. Nonetheless, they are instance diagrams conforming to an underlying UML model. This is quite powerful from a software engineering point of view. First, the layered diagrams can be used in customer meetings for requirements elicitation and refinement. Second, they can be used as a basis for code generation.
3) **Layered Set.** Given the recursive definition of the ecosystems arena of operations and the RoIs, the ZEDs are a layered collection of diagrams, and certain concepts hold and are valuable. The concourse in- and outflows between Layers 1 and 2 (see Figure 9 and Figure 10) are balanced and should remain so. In both diagrams, there is one ingress and two egress flows. Similarly, the input and output attributes and credentials are also balanced from level to level.

3.5 Chapter Close

It is helpful to keep in mind the following as we proceed with credential design, distributed processing, and risk-aware access control:

1) **Cyber-Physical Ecosystem.** A transactional exchange protocol in which data elements are assured by cryptographic credentials helps the ecosystem to flow smoothly by protecting the interests of the stakeholders.
2) **General Credential Perspective.** The credential is a secure data transfer object that encloses signed attribute attestations. Credentials are signed by the issuer, held by the subject, and verified by downstream service providers. The credential provides a means to satisfy the data requirements of all parties: security (including non-transferability), non-repudiation, verifiability, and reliability.
3) **Checkpoints.** A checkpoint is situated in the cyber-physical terrain and acts as a control point through which a subject must apply for passage. The checkpoint consumes electronic data, verifies their integrity, conducts risk analyses, and routes the subject accordingly, granting or denying their request. The checkpoint may emit digital data or credentials as transaction outputs.
4) **Stakeholders.** There are various stakeholders – subjects and service providers (or “authorities”) – in the system. From a credential perspective, the subject is the credential applicant, recipient, and redeemer, while the service providers are credential issuers, verifiers, or both. Stakeholders have independent goals that may be in conflict when examined in the context of either the stakeholders themselves, the stakeholders and their peers, or the stakeholders and the

ecosystem-wide collection of objectives. These conflicting goals can, at times, place the stakeholders at odds with one another.

Chapter 4 Credential Design 1: Cryptographic Envelope

4.1 Chapter Introduction

This chapter presents Credential $Design\ 1$ – an envelope-based design for the Mobile Travel Credential (mTA). $Design\ 1$ uses fuzzy extractors for key generation, symmetric encryption, and public key cryptography (asymmetric encryption and signature scheme) for envelope security. $Design\ 1$ allows the issuer and verifier to use distinct biometric matchers with dynamically configurable match thresholds. As will be seen in the property analysis, $Design\ 1$ has the advantage of ease of adoption by the issuing and verifying authorities at the expense of some privacy features. $Design\ 1$ was produced as part of a series of designs and projects with the Canadian Government notably, the Canada Border Services Agency (CBSA), the Department of Immigration Refugee, and Citizenship Canada (IRCC) and Defense Research and Development Canada (DRDC). The designs were produced independently by the author as part of this thesis, and implemented with a project team from the University of Ottawa.

4.1.1 Scope

4.1.1.1 Publications

Material and results from this and chapter 5 have been published in (Bissessar et al., 2014)(Bissessar et al., 2018)(DRDC, 2013)(DRDC, 2017)(DRDC, 2018).

4.1.1.2 Contributions

The contributions of this chapter are as follows:

1) **A Non-lendable EbC design (**$Design\ 1$)

An envelope-based design for the Mobile Travel Credential (mTA) is provided using standard cryptography, commercial biometric matchers, and fuzzy extractors for key generation. The cleartext, signed EbC is a common credential design used in DTC and mDL (ICAO 2020)(ISO, 18013)**.** To our knowledge, this is the first illustration of fuzzy extractors and envelope encryption as addition to the EbC and its assessment against security, privacy and operational properties.

.

2) **A non-lendable AbC design (**$Design\ 2$)

An attribute-based design for the Mobile Travel Credential (mTA) is provided using Brands' Digital Credentials (Brands, 2000)and fuzzy extractors for non-lendability. The non-lendability design is based on (Bissessar et al., 2014) adapting it to use fuzzy extractors for face. This can also be used on Anonymous Credential. Other approaches to non-lendability exist and will continue to be developed(Sarier, 2021). To our knowledge, this is the first application of fuzzy extractors for face demonstrated using AbC on the border and travel domain.

3) **Assessment of EbC and AbC designs against target properties**

Chapters 4,5 and 7 assess the EbC and AbC independently and with respect to each other against a framework of security, privacy and operational properties. While different AbC schemes have been compared for functionality and performance in a number of studies (c.f. (Veseli and Serna, 2016)(Baldimtsi and

Lysyanskaya, 2013b)), we have not seen any work contrasting EbC and AbC designs against a set of engineering properties. We believe this is a valuable contribution to engineers and applied scientists in the field. To our knowledge, this is the first time a set of properties including security privacy and operational aspects are used to assess these important alternative design patterns comparatively. The operational assessment makes evident challenges to AbC and fuzzy extractors deployment in an ecosystem, including ease-of-adoption, interoperability, performance measurement and adaptability. The operational properties assessment uncovers some important areas of required work for making fuzzy extractors operationally usable, including ROC Curve analysis, and the ability to decouple issuer and verifier error thresholds. The identification of these operational challenges is novel in the fuzzy extractor literature and pose these as open problems.

4.1.1.3 Exclusions

1) **Comparative Fuzzy Extractor Performance Metrics.** This thesis does not present the transaction data the fuzzy extractors or the commercial matchers. The fidelity of the commercial systems vs. the fuzzy extractor was observed in field experiment reports but data were not collected (DRDC, 2018).

4.1.2 Target Properties for Design 1

Credential *Design 1* focuses on ease of adoption, interoperability, and security with some biometric privacy. We propose a credential design based on EbC that

is readily adoptable in enterprises today. An EbC does not include the advanced privacy features of AbC, which are featured in *Design 2*. Non-lendability and biometric privacy are of primary concern. *Design 1* uses fuzzy extractors for key generation, and biometric matchers for face verification. Table 6 lists the target properties of Chapter 1 and prioritizes them for *Design 1*. The platform security privacy respect assumption on the Smartphone is removed in the discussion *Design 1*.

Table 6) Target Properties for Design 1

Category	Property	Priority
Security	Unforgeability	High Priority
	Tamper resistance	High Priority
	Non-transferability	High Priority
Privacy	Unlinkability	Reduced Priority
	Composability	Reduced priority
	Selective-show	Reduced priority
	Biometric privacy	High Priority
Operational	Interoperability	High Priority
	Classifier accuracy	High Priority
	Data usability	High Priority
	Adaptability	High Priority

Security properties are mandatory. It should not be possible for an attacker to create a bogus credential by forging the signature of bona fide issuers. The credential package must be tamper-resistant. It may be assumed that the subject carries required credential data and that the subject does not need to see the attribute values within the CC. It should not be possible for parties outside the immediate transaction to see any biometric information at rest or over the air. There should be no storage of biometric images – transacting authorities should not store biometrics. Biometric performance and credential-level attributes may be achieved at the expense of composability.

In terms of operational properties, verifiers should be able to perform meaningful conformance and risk analysis of user-supplied credential data. Interoperability is a high priority. An approach is needed that can receive high uptake across stakeholders. One credential should be reusable across the issuance and multiple invocations of the verification protocol without leaving a linkable trace on the public transcript.

4.2 Building Blocks

4.2.1 Secure Sketches

The secure sketch defines two functions $\langle sketch,\ rec \rangle$. This pair of functions allows a secret w from a metric space to be securely "sketched" into public data p, which can be

safely stored and later used to recover w if a sufficiently similar message w' is provided, where $p = sketch(w)\ and\ w = rec(w', p)$. In the context of our discussion, the secret w and the sufficiently similar w' are biometric templates. The fuzzy extractor is defined by the functions $\langle gen,\ rep \rangle$, which allow a key k to be generated and reproduced from a biometric w and a sufficiently similar w'. As with the secure sketch, the fuzzy extractor generates auxiliary helper data, which must be provided at the time of key regeneration. A fuzzy extractor is a pair of algorithms $(gen(\ldots), rep(\ldots))$ that allow randomness to be extracted from an input string and later reproduced exactly using another input string sufficiently close to the original. Most fuzzy extractor schemes also produce helper data, which are created during the initial generation step and used to assist in reproducing the randomness. The pair of algorithms can be represented as in Figure 10.

$< P, R > = gen(b)\ and\ R\ =\ rep(b', P)$, where:

P denotes public data that are safe for storage and used to assist in the $rep(\ldots)$ algorithm;

R is a random string that can be used for cryptographic purposes;

b is an input string, for example a biometric template; and

$b'$is fresh input within a certain similarity distance t from b.

Figure 11) Fuzzy Extractor Methods

4.2.2 Biometric Matchers

A biometric verification matcher can be represented as object M with operations $M.train(G)\ and\ M.verify(b_0, b_1, t)\ subject\ to\ e$, which initializes M and then applies it on candidate biometrics at a given threshold. M.verify() returns a Boolean – whether or not $b_0,\ b_1$ are classified as similar, within distance t, subject to an error term that captures the expected rates of Type 1 and Type 2 errors. To provide an example of how

these methods can be implemented in terms of a classifier or face matcher, consider the eigenfaces algorithm: $train(G)$ is defined to initialize the feature extractor-based on the images in the training gallery G. In the case of PCA, the feature extractor is a matrix that verifies and then uses the feature extractor to create two lower-dimensional templates t_0, t_1 from the input face images b_0, b_1 and returns true if templates t_0, t_1 are within Euclidean distance t of each other. In $Design\ 1$, the issuer and verifier may have distinct biometric matchers M_I and M_V.

4.3 Proposal

For the purposes of Chapters 3 and 4, we focus on the mTA issuance and verification. Furthermore, to discuss the verification algorithm, we concentrate on the border services authority. As will be discussed in the analysis sections, the issuance algorithm would work equally well if the target credential were a notional “virtual boarding pass” or a “virtual custom stamp.” Similarly, the verification algorithm applies whether the verifier is the airline or the border services agency. Check-in kiosks, for example, may be used by airlines or airport authorities prior to departure or by border arrival authorities at the destination.

4.3.1 Issuance

4.3.1.1 Algorithm Description

The credential “Issuance” process is a request-response protocol between the user system p_i, and the issuance service p_I.

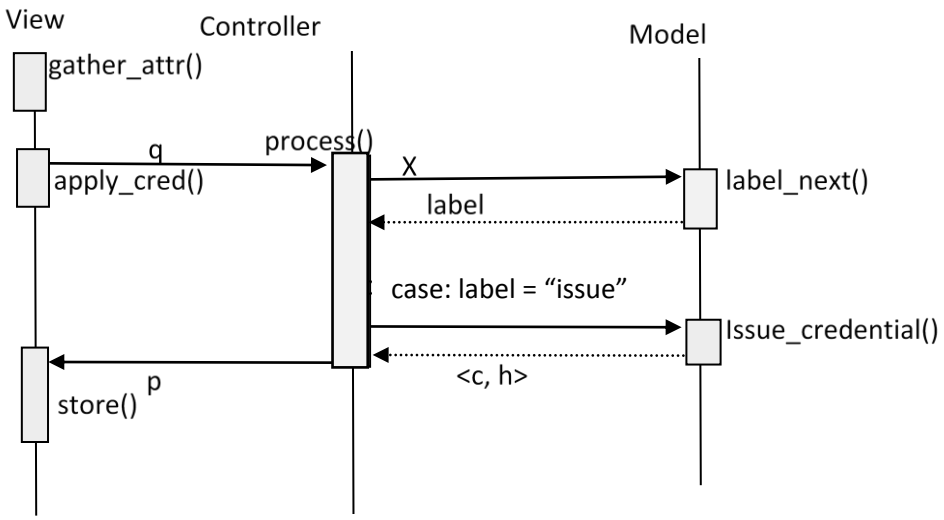

Figure 12) Generic Credential Issuance Protocol

The user agent (UA) acquires travel information X, enrolment biometric $\mathrm{b_e}$ and passport data D $\langle X,\ \mathrm{b_e}, \mathrm{D}\rangle$ and prepares them as request q, to be sent to a_1 for processing. The cloud-based system of $a_1$receives q and sequences the core algorithms of the compliance and risk modules $\mathcal{R}_I = \{\mathcal{R}_{I_1}, \ldots, \mathcal{R}_{I_n}\}$, and the credential completion algorithm do_issue() to create credential c. On successful completion of the issuance protocol, $a_1$packages credential c into response r, which is transferred to s_i for appropriate storage.

Based on the assigned label, I conducts the appropriate follow-up processing, which could be to refuse the application or to escalate the case to supporting systems (which may include manual processes). If assessment yields acceptance, then the response c

is created. This credential is a cryptographic object, signed by the issuer and encrypted for the verifier.

$Design\ 1$ uses the fuzzy extractor as follows. The traveler's facial biometrics at the time of issuance (i.e., the selfie photo) are input to the fuzzy extractor to produce an AES key and the helper data. The AES key is used to encrypt the credential, and the helper data can be stored on the phone along with the credential. The AES key itself is encrypted using the public key of the kiosk at the destination airport. In our system, this derived AES key is referred to as a "renewable biometric reference" (RBR).

Note that since communication between the mobile device and the issuance server occurs over the Internet, a secure communication channel must be established between this client and server. The obvious choice would be to use TLS so that the server can be cryptographically authenticated, and the mobile device can transfer the relevant travel application data with confidence.

4.3.1.2 Communications

HTTPS is used to establish a secure communication channel for all communicating entities (Rescorla, 2001)(Sherif, 2016). HTTPS uses TLS to provide a secure channel for the exchange of data by implementing encryption and certificate-based authentication.

4.3.1.3 Key Generation

Key generation algorithms are as per the ECDSA (Johnson et al., 2001), RSA (Rivest et al., 1978), and AES (Daemen and Rijmen, 2002). Storage uses X.509 encoding for public keys and PKCS-8 encoding for private keys.

4.3.1.4 Credential Issuance

The credential was signed using the issuer's private ECDSA key. Thereafter, the signed credential was encrypted using RBR as an AES key, and RBR was encrypted using the RSA public key of the kiosk. This whole package was then stored on the traveler's phone.

4.3.1.5 Issuance Data Acquisition

The answers to the questionnaire X_I were obtained from keyboard input. Facial biometric data b_{e} were obtained using the smartphone's front-facing camera. The application rendered an oval overlay to help guide the user through photo capture, and open-source optical character recognition was used to extract the keying fields needed by the Basic Access Control (BAC) protocol between chip and phone from the machine readable zone mrz on the e-passport biodata page. Reading D from m was achieved using JMRTD (JMRTD, 2018), which also provided facilities to parse D and complete ICAO passive authentication.

4.3.1.6 Issuance Risk Assessment and Conformance

Biometric risk was simplified to threshold match scores of a commercial 1:1 face recognition engine. Thus, $\mathcal{R}_{I_b} = B_I(\mathcal{M}_I, t_I, b_o, b_{\mathrm{e}})$. Furthermore, document risk was

assessed in two manners: first, as a function of the BAC and ICAO integrity checks on D and second, as a simulated call to the Interpol “Stolen and Lost Travel Document” database (Safjanski, 2015). Finally, attribute risk was illustrated using mock business rules to assign country-specific risk given the input data X_I.

4.3.1.7 Prototype-RBR Generation

Having passed risk and conformance assessment, the last step was to generate the credential itself. The biometrically derived AES key (RBR) was created using the fuzzy extractor *gen(…)* function.

4.3.2 Verification

4.3.2.1 Algorithm Description

As above, we assume that the global environments and subject and verifier computers p_s and p_V have been properly initialized. As such, we assume that p_v has a sensor capable of capturing facial images suitable for biometric comparison, near-field communication for passport and kiosk interface, data connectivity within the ecosystem, processing power, and trusted secure storage. Moreover, we assume that p_V has access to core systems and external Oracles, particularly matcher M and directory D. Verifier v has previously deployed the verification point of presence p_V within a designated context in the application process on the UA.

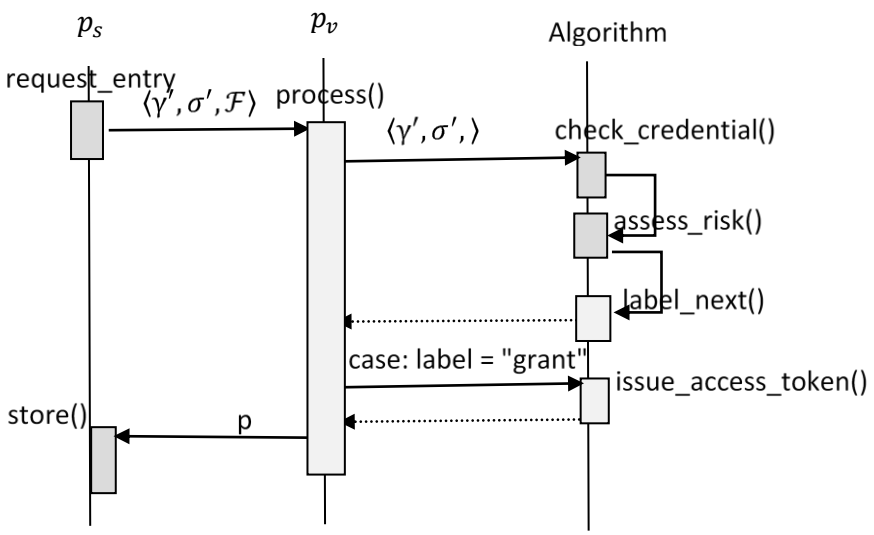

Figure 13) Generic Verification Protocol

Figure 13 illustrates the workflow between the smartphone and the kiosk, highlighting the key algorithms at the time of verification. The kiosk application guides the user through the sequence, instructing the traveler to put their passport on the document reader and steering the traveler through the acquisition process for the verification-time facial image. Once the passport is read and the fresh image is captured, the mobile credential from the phone is transferred via NFC to the kiosk. At this point, the kiosk uses a commercial matcher $\mathcal{M}_V$ to perform biometric matching between the chip image and the fresh facial image. If the match is within the accepted threshold for the border security process, the kiosk verifies the credential obtained from the phone. If the credential is valid, the next screen shows that the application was successful; otherwise, it shows an error code, and the traveler is referred to the border service officer to solve the problem.

Upon arrival, the traveler's photo taken by the kiosk is used to re-derive the AES key (along with the helper data downloaded from the phone). This key is then compared with the AES key decrypted using the kiosk's private key (alternatively, the re-derived key is used to attempt to decrypt the credential that has been downloaded from the phone). If the kiosk photo and helper data can generate the correct AES key, then the traveler in question is highly likely to be the person who applied for the credential at the time of issuance. This RBR technology allows for biometric verification without the use of a stored biometric template.

In this verification process, the kiosk can confirm the binding between the human user and the passport (using accurate face-matching technology), between the human user and the mobile travel credential (using a privacy-respecting biometric approach),

and between the passport and the mobile travel credential (using digital signature technology over the data from the passport that are also encoded in the credential). This conjunction of bindings enhances security over the processes in place today, wherein a human officer does a manual face match between the traveler standing in front of them and the image on the passport bio-data page.

4.3.2.2 Credential Verification

Verifier V commissions self-service kiosk K, which is responsible for presenting a questionnaire Q_v, processing the answers, capturing the verification biometrics, reading the passport, assessing risk, and labeling the outcome (which can include manual intervention or automated approval).

4.3.2.3 Verification: Data Acquisition

As represented in Figure 14, standard kiosk features are used to acquire X_v mrz, P_{bd}, and D. Here, X_v is obtained by user input on the soft keyboard of K, and the biometric b_v of U is captured by the kiosk's image capture camera. The document reader of k captures an image of P_{bd}, obtains mrz, and uses it to read D.

$X_v = k.kb(Q_v)$ $b_v = k.c_1(U)$ $mrz = k.doc\ (P_{bd})$ D = m.nfc(mrz)
Figure 14) Data Acquisition at the Time of Verification

In the prototype, k did not have facilities to read c. A custom prototype peripheral was hence constructed to do so. This module consisted of a cradle to hold the phone

against an nfc reader (ACS, 2018), a camera, and an ODroid controller (XU4, 2018). The traveler inserted m into the cradle, at which point the kiosk's NFC module could download c.

4.3.2.4 Credential Opening

Having obtained c from the NFC module, the kiosk first verifies the signed credential against the issuer verification key Ipk and then applies a two-step decryption process (see Figure 15). First, Vsk is used with $rsa.dec_{Vsk}$ to decrypt c_1. This yields symmetric key k_s, which is then used to decrypt c_2 with $aes.dec$ to produce $< Y >$.

if $ecdsa.ver_{Ipk}(c,\ \sigma_I)$: $\mathrm{k}_s = rsa.dec_{Vsk}(\mathrm{c}_1)$ $< Y > = aes.dec_{\mathrm{k}_s}(\mathrm{c}_2)$
Figure 15) Credential Opening

4.3.2.5 Conformance and Risk Assessment

The sequence of validations completing the verification protocol are as follows:

1) The freshly obtained biometric is compared against the extracted passport photo. A favorable match suggests that the individual at the kiosk is the rightful passport holder.
2) The rbr derived from the kiosk photo (and the helper data from the phone) is used to decrypt the downloaded credential. Success indicates that the credential holder (physically present at the kiosk) is the same individual to whom the mobile travel credential was originally issued.

3) The expiration date of the credential mtc_exp_dt must be greater than the current date. This ensures that travel is within the permitted mTA period.
4) An online call to the issuer is made to verify that the mTA is not revoked. This ensures that no recent problems exist in the current subject's case.

These checks collectively ensure the entitlement–ownership relationship: the subject at the kiosk is the passport holder; the passport holder is the person to whom the mTA was issued; and the mTA is still valid.

4.4 Properties Analysis

The credential for $Design\ 1$ is structured as depicted in Figure 16. This structure can be viewed in three layers: c, c_1, and c_2. The credential itself is represented as $c = (c_2, \sigma)$ in which σ is the issuer's signature on ciphertext c_1 (a DSA signature produced using the issuer's private signing key $S_{a1\_k\_pr}$), and c_1 is an RSA encryption of m_1, produced by the issuer using the public key $E_{a2\_k\_pu}$ of the intended verifier a_2. Here, m_1 is the tuple (c_2, h, r) in which c_2 encrypts the required attributes using the symmetric key E_{u_bdk} generated using the fuzzy extractor, and h, r are the associated helper data and randomness used in the fuzzy extractor.

$c = (c_2, \sigma)$ s.t. $\left(\sigma := (c_1 := (\ m_1 = (c_2 := (m)_{E_{u\_bdk}},\ h, r\ ))_{E_{a2\_k\_pu}}\right)_{S_{a1\_k\_pr}}$
Figure 16) Design 1 Credential Structure

Furthermore, c_2 protects attribute data, under a biometrically derived key that is under the subject's control, and c_1 includes data required for the regeneration of the biometrically derived key. In addition, σ is a signature on the entire envelope, and r can be input separately for added security, at a small cost in terms of user convenience, storage, and communications.

The key $E_{u_{bdk}}$ is discarded after c_2 is created and must be recreated during the verification process. Public asymmetric keys $E_{a2\_k\_pu}$ and $V_{a1\_k\_pu}$ are retrieved from the ecosystem directory services.

4.4.1 Security Properties

$Design\ 1$ features the security properties of unforgeability, tamper resistance, and nontransferability. Unforgeability and tamper resistance are delivered by the DSA algorithm, and nontransferability builds on these and the similarity properties of fuzzy extractors.

The credential consists of DSA signature pair (c_1, σ); the tamper resistance of the DSA protects the credential's integrity. Any change $c_1{}'$ would be detected $DSA.ver_{a1kpu}(c_1{}', \sigma) == fail$, and the credential is thus tamper-resistant.

The credential is also not forgeable. Assume that the attacker seeks to make a credential on falsified data on $\mathrm{p2}'$. To do this, the attacker would need a private signature key $\widetilde{a1}kpr$ such that $\tilde{\sigma} = DSA.sign\_\widetilde{a1}kpr(c_1{}')$ and $DSA.ver(c_1{}', \sigma) == 1$. However, the DSA is secure against the derivation of the private key given the public data; therefore, forging the signature $\tilde{\sigma}$ would require the attacker to break the DSA. It is also possible to forge a signature by finding a collision in the hash function, but this can be mitigated if a collision-resistant cryptographic hash function is used in the signing process.

Furthermore, nontransferability is achieved through the properties of the fuzzy extractor and by virtue of tamper resistance. The fuzzy extractor gen and rec methods are defined over the metric space over which an error-correcting code is defined. The key created by the fuzzy extractor gen method can only be recreated by a biometric that is within the designed error correcting distance of the original. FE protects against lending precisely because the biometrics from the borrower are assumed to be outside the error-correction tolerance. The biometrics of the imposter applied to the fuzzy

extractor rec method would create the incorrect bit string, say r. $AES_r.decrypt\ (c)\ != m$. As such, no useable data would be retrieved for the attacker, and the credential would thus not be useable.

4.4.2 Privacy Properties

$Design\ 1$ features the privacy properties of biometric privacy, but not of unlinkability, composability, or selective show.

$Design\ 1$ delivers biometric privacy by two features. First, no biometric data are embedded in the credential. Second, there is an assumption of honest-but-curious transactors, who agree to delete the biometric image after its sanctioned use. The biometric key is derived from source biometric b_e after it is compared with the passport picture b_0. Neither b_0 nor b_e are stored in the credential; they are, however, used by the issuer to verify the biometric identity of the subject and to quantify biometric risk. Similarly, at the time of verification, b_0 and b_e are obtained from the subject and used by the authority to verify the subject's identity and quantify biometric risk. After b_v is used (along with h and r) to regenerate symmetric key u_{bdk}, all biometric data are deleted. Moreover, u_{bdk} itself is also not stored in the credential.

$Design\ 1$ does not provide unlinkability, because the subject does not obfuscate the credential (c_2, σ) between the issue and verification transactions. As such, the audit trails of both the issue and verification transactions on the transcript include (c_2, σ), making the transactions linkable.

$Design\ 1$ does not provide composability or selective show, since the credential attributes within m cannot be individually separated and offered by the subject to

compose subset proofs or linked proofs across credentials. The design of CD_1 is highly coupled to the specific needs of the verifier.

4.4.3 Operational Properties

$Design\ 1$ provides the operational properties of interoperability, prediction accuracy, and adaptability. As discussed, authorities require performant transactions, which provide meaningful data and adaptable algorithms to best react to risk and opportunity in operations.

Our focus is on mitigating the risk of identity fraud using biometrics. We discuss a) the three properties in the context of biometrics and risk and b) the design choice of using classifier-based matching or fuzzy extractor matching.

In terms of interoperability, we chose to use classifier-based matching for biometric risk, and fuzzy extractors for key generation. Choosing commercial matchers imposes a requirement on the interface between the subject and authority systems: biometric images must be sent to the authority for comparison. Having imposed that requirement, however, the authority is free to use any number of fit-for-purpose matchers and is decoupled from decisions that other authorities may make. Our use of fuzzy extractors occurs after biometric risk has been addressed by the classifiers. Once the subject has been suitably identified, the fuzzy extractor keygen is performed. These are two separate operations.

With the use of fuzzy extractors, common practice is to rely on their success or failure in key gen or regen as a test of biometric identity. While this is a conceivable approach for our problem, we chose not to pursue it because with fuzzy extractors,

interoperability and adaptability are more difficult. The main reason for this is that the fuzzy extractor error-correcting code parameters must be propagated across the ecosystem to be common to issuers and verifiers, which is significantly more taxing on interoperability than simply passing standardized image formats.

The choice of biometric matcher versus fuzzy extractor also highlights the property of prediction correctness. By choosing a biometric matcher, we allow authorities to select the most appropriate classifier for their operations – this may be a commercial matcher, a cloud service, or an internal design of ML/AI components. The fuzzy extractor design, on the other hand, relies on the error-correcting code approach as deployed across the ecosystem. The negative characteristics experienced for interoperability also hold for prediction accuracy. Therefore, the choice of a biometric matcher is superior for prediction accuracy as well.

Finally, and perhaps most importantly for our thesis and its direction towards continuous risk-balanced access control, there is the operational property of adaptability. Authorities must be free to calibrate their systems to the demands of their operational contexts, which include challenges due to the elements (e.g., lighting, temperature, motion blur, etc.) and daily actualities (e.g., amber alerts, high volume predictions, rerouted flights). Faced with this need, classifier-based technology allows our authorities to set operational thresholds sensitive to the perceived risk.

The difficulties of fuzzy extractors in this environment follow a discussion similar to the one above, on the (m, l, e) parameters. Since k and h are produced according to an ECC parameter that must be common across enrolment and verification, the verifier does not have the flexibility to change a threshold to alter the rates of error. For reasons

of verifier-configurable thresholds, the classifier approach offers higher adaptability than the fuzzy extractor approach.

4.4.4 Variations on Design 1

$Design\ 1$ and variations of it were implemented in a prototype environment (Bissessar et al., 2018). This section discusses some of the associated drawbacks and modifications.

4.4.4.1 Key Release

The interplay between thresholds for biometric matching and key generation are discussed in Section 4.4.1. In a key release approach (Cavoukian and Stoianov 2009), rather than using a fuzzy extractor to generate a biometrically derived key, a standard key is used. The key is conditionally "released" either by a secure element smartcard or the application logic in the verification protocol on a successful match.

4.4.4.2 Broadcast Encryption

The fact that the algorithm encrypts the package with $a2\_k\_pu$ also presents a limitation. Recall that the airline and the border both perform a verification on the mTA. Under the baseline approach for D_1, the credential is encrypted for one recipient. This can be remedied by using a broadcast encryption scheme such as (Fiat and Naor, 1994). Adopting such an approach would have advantages. For instance, the encryption could be sent and used by multiple verifiers. This approach would also support the revocation of read permissions for verifiers. However, a broadcast encryption approach does not meet the requirement that standard cryptography be used with ease for IT security departments.

4.4.4.3 Transparent Envelope

The problem illustrated in Section 4.4.4.2 is one reason that, in common practice, envelope-based credential designs typically do not encrypt the attributes. The ICAO digital travel credential (ICAO 2020) and the ISO mobile driver's license (ISO, 18013) are examples of envelope-based designs that do not encrypt attribute data, but merely sign the envelope.

4.5 Chapter Summary

Design 1 proposes an envelope-based design in which all attributes required by the verifier are signed and encrypted by the issuer. The holder stores, forwards, and sends the credential, as received, to the verifier.

The use of standard digital signatures and credentials allows for ease of adoption, and interoperability is a motivating factor for stakeholders. Biometric performance and fine control over biometric error rates are required by operational agencies. As such, fit-for-purpose matchers (possibly different between issuer and verifiers) are employed.

The EbC design of *Design 1* does not support the data minimization or composition features of AbC, which are featured in $Chapter\ 5$. However, nontransferability and biometric privacy are of primary concern to citizens and authorities. The mismatch between the acceptance rates of biometric matchers and fuzzy extractors has been experienced in the field. A solution to this problem is a key-release approach in which a cryptographic key is released upon successful matching of a custom matcher.

Chapter 5 Credential Design 2: Attribute-based Credentials

5.1 *Chapter Introduction*

This chapter proposes an alternative credential design (*Design 2)* that uses Brands' attribute-based credentials. *Design 2* addresses the privacy requirements of selective show and credential composition, and non-lendability and biometric privacy are achieved using fuzzy extractors. The use of attribute-based credentials in *Design 2* yields privacy and security advantages but also presents some drawbacks in terms ecosystem standardization and adaptability. *Design 1* and *Design 2* are compared in Chapter 7.

5.1.1 Target Properties for Design 2

Table 7 presents the target properties for *Design 2*. Depending on the ecosystem development methodology (for example from an incremental prototyping perspective), it may be difficult to lower the bar from the successes achieved by *Design 1*, which achieved "security" and "operational" qualities in Ch.4. As such. becomes difficult any alternative to sacrifice some of these qualities. As such the success criterion for the AbC design is somewhat higher: each of the privacy requirements and security should be addressed, and operational properties must be maintained. As we shall see, by the end of this chapter, not all desired properties are equally achievable. Adoption of any

alternative based on EbC or AbC becomes a matter of the collective objectives of the ecosystem and stakeholder, which is highlights the motivation and need for credential design as an area for engineering focus.

Table 7) Target Properties for Design 2

Category	Property	Priority
Security	Unforgeability	High Priority
	Tamper resistance	High Priority
	Non-transferability	High Priority
Privacy	Unlinkability	High Priority
	Composability	High priority
	Selective-show	High priority
	Biometric privacy	High Priority
Operational	Interoperability	High Priority
	Classifier accuracy	High Priority
	Data usability	High Priority
	Adaptability	High Priority

5.2 Building Blocks

5.2.1 Attribute-based Credentials

Various credential systems have been proposed in the literature, though the two predominant systems today are Digital Credentials and Anonymous Credentials.

Credential systems generally include three entities: the individual, the issuer, and the verifier. The individual applies for a credential from the issuer by submitting attributes. The individual receives a signed data package that is later shown to a verifier to claim a privilege.

Anonymous credentials and digital credentials have certain similarities. The general protocols both include Issue, Show, and Blinding operations. In the Issue protocols of both schemes, User U sends attributes X to issuing organization I, which creates signed credential C, which is returned to U. Both schemes provide a blinding operation which U applies to the signature and credential prior to the show protocol to prevent transaction tracing. Both schemes also include a show protocol in which U presents the signed credential to a verifier V, and U makes a provable claim involving attributes of X. However, the schemes also differ in some important ways. In the show protocol of anonymous credentials, U shows neither the original credential nor the actual values of the attributes to the verifier. Rather, blinded credential, signature, and attributes are sent, and a Zero Knowledge Proof of Knowledge (ZKPoK) is used to convince the verifier of accuracy of the credential, signature, and attribute values. This provides additional privacy at incremental computation costs. This chapter outlines the protocol for digital credentials; however, anonymous credentials can also be used.

5.2.2 Pedersen Commitments

The Pedersen commitment allows senders to create a publicly storable commitment on a value that binds to the value and perfectly hides the value from being derived.

The Pedersen commitment scheme has two protocols $C_s = Commit(s,r) = g^s h^r \ (mod\ p)$ and $true|false = Open\ (C_s, s, r)$, where the secret s is a value from Z_q, and random value *r* is uniformly drawn from Z_q. The specification of $(mod\ p)$ for the Pedersen commitments are not included in the present chapter but are implied by context. Here, the commitment is on the biometrically derived cryptographic key: the "hiding" property preserves the privacy of the key, and the "binding" property ensures security.

5.2.3 Zero Knowledge Proofs of Knowledge

A ZKPoK is an interactive protocol in which a prover P convinces a verifier V of possession of knowledge, without divulging that knowledge. In general, a Proof of Knowledge (PoK) has the characteristics of completeness and validity. The property of completeness means that, if P holds the required knowledge, P will succeed in convincing the verifier V of this. The property of validity means that, if verifier *V* accepts the proof, P has the required knowledge. An additional property, zero-knowledge, can be added; this means that, during the protocol, V learns nothing other than that P holds the required knowledge. A PoK with completeness, validity, and zero-knowledge is called a ZKPoK.

5.3 Proposal

$Design\ 2$ uses a fuzzy extractor and attribute-based credentials construction. This section introduces the main entities involved in the scenario, describes the traveller flows, and describes the algorithms that are significant from the perspectives of security

and privacy. In addition, certain special features of the algorithm, including the ability to help detect passport fraud at pre-board (optional), triangular binding of mTA to a portrait-face RBR, and e-passport scalability to a paperless “m-passport” application (with the potential for the e-passport to be optional after the first verification), are outlined.

5.3.1 Keygen at Issuance

Travelers’ smartphone captures facial photo to generate the RBR that is subsequently submitted to the issuer to be sealed into the mTA.

Figure 17 demonstrates the creation of a Pederson Commitment C_I using facial image b_I and random data r_1 and r_2. Arguments b_I and r_1 are are used to make the fuzzy extractor output tuple (k, p) consisting of biometrically derived key k, and public data p. A Pederson commitment is then made on k and r_1 on the reserved Digital Credential generator bases say $g_k, and\ g_r$. The commitment C_I is sent to the issuer along with the traveller’s other attributes and issuance proceeds as per the Brands protocol. The subject retains rbr, k, p, r_1, r_2 for use in the verification protocol.

$SE :: make\_rbr(\ b_I, r_1, r_2)$:

$(k, p) = FE :: gen(b, r1)$

$C_I = PC :: Commit(k, r2)$

$return\ (\ C_I, k, p)$

Figure 17) Creating RBR on Secure Element

5.3.2 Generation of Verification RBR

On arrival at the destination airport, a traveler's ownership of the mTA must be established. To do so, the kiosk creates a verification-time RBR C_S. To create C_S, the traveller initiates communication with the kiosk, transferring fuzzy extractor public data p, and fresh randonmness r_3. A fresh biometric b_v is captured by the kiosk. As shown in Figure 18, biometric b_v and public helper data p are passed to FE::$rep(\dots)$ which regenerates a biometric key k'. If $b_I \; and \; b_v$ are sufficiently similar, the regenerated key k' will be the same as the original k_I which had been generated in the issuance protocol.

$SE :: rbr\_regen(\, b_v, p, r_3)$:

$k' \; = \; FE :: rep(b, p)$

$C_S = \; PC :: Commit(k', r_3)$

$return \; (\, C_S, k')$

Figure 18) RBR Creation at the Time of Verification

The fresh Pedersen commitment C_S is created on attribute data k' and r_3 as discrete log representations on public generator bases $g_k, and \; g_r$. C_S and k' are provided to the traveler for use in subsequent steps in the verification process.

5.3.3 Algorithm: Show Protocol

After the RBR has been regenerated by the arrival kiosk, the final step in the workflow requires verification of the mTA, the e-passport, ownership of them, and any claimed travel privilege. This requires the verifier to be sure that:

1) The mTA data package has not been tampered with,
2) The RBR sealed into the mTA and regenerated in the previous step are for on a biometrically similar photo,
3) The e-passport number within the mTA corresponds to that of the passport held by the traveller,
4) The traveler's face and the photo on the e-passport match,
5) The traveller claim of privilege is valid.

To verify that the digital package has not been tampered with, a verification relation is evaluated by the verifier. This verification relation is defined by the underlying credential scheme. In general, this is a function of the credential itself and the issuer's public key. The public key may be installed on the kiosk itself or accessed through an online connection.

Once the verification relation has been checked, the traveller must prove ownership of mTA by showing that C_S (the RBR created at the time of verification) and C_I (the RBR within the mTA) are on the same derived key k, and thus derived from biometric samples belonging to the same individual. The commitment generated at verification time, can be signed by the individual and may serve then, as a type of authentication credential[2].

[2] The scheme used in (Bissessar et al., 2014) has the interesting property that the commitment issued at verification can be made into a credential by signing it. This can provide the subject to produce authentication credentials at will.

Following proof of biometric ownership, the smartphone and kiosk perform compliance checks on the credential data – for example that the mTA allows entry into the country, that the entry dates are appropriate, and that the mTA has not been revoked. This is conducted using a combination of the statement proof mechanism of the underlying credential system and online revocation checks with the issuing authority.

5.4 Properties Analysis

Security Properties. $Design\ 2$ security properties of unforgeability and tamper-resistance are the same as those proved by (Brands, 2000). *Design 2* uses a fuzzy extractor-based design with an assumption of verifier honesty. The assumption of verifier honesty can be relaxed if the fuzzy extractor is moved to a secure platform on the smartphone.

Privacy Properties. *Design 2* features single show unlinkability at the cryptographic primitive level and biometric privacy selective show. In the Brands digital credential scheme, the blinding function transforms the issued credential signature pair used in verification such that they cannot be correlated in single-verification use. The blinding function in Brands digital credentials does not allow the credential to be used multiple times. The credential, signature pair are traceable if used in multiple verification protocols. Further, in the Brands digital credential scheme, attributes are sent as clear text to the verifier. For this reason, an attacker could be able to correlate transactions based on the attributes sent, though conducting the Brands protocol using HTTPS could solve this problem. If multiple show unlinkability is required at the cryptographic primitive level, the protocol outlined by (Camenisch and Lysyanskaya, 2001) could be applied.

Design 2 offers biometric privacy, which relies on the honesty of the verifier engaged in a transaction. The biometric images are sent to the verifier, and it is expected that they are deleted after use. A corrupt verifier may ignore the convention of deleting images and save biometric information. Modifying the design offers a solution to this problem. In *Design 2*, the fuzzy extractor ***rep(…)*** occurs on the kiosk; however, moving it to the user device achieves stronger biometric privacy. The assumption of a trusted user device and client application is increasingly important. In today's environment of "Bring Your Own Device" (BYOD), user devices are rife with trojans and data hungry adware -- such an assumption is not realistic. *Design 2* allows for composability and selective show. In both Brands and Camenisch Lysyanskaya credentials, attributes from two different credentials may be combined. These features allow data to be combined across credentials, while only the required attributes are sent to the verifier. These features are essential in a privacy-by-design approach.

Operational Properties. *Design 2,* which is based on Brands digital credentials, can support the intelligence function of human and computer analytics. However, the issuing and verifying authorities, must adopt non-standard cryptography and must accept the incremental complexity and uncertainty of developing around a non-standard set of cryptographic protocols. Interoperability is challenging, as all service providers are tied to the error correction parameters for fuzzy extractor key generation. In addition, the Brands digital credential scheme imposes global parameters on the environment.

Data Reliability. *Design 2* offers data reliability due to the security properties of non-transferability, unforgeability, tamper resistance of the credential, and the use of actual attribute values or optimized representations in support service requests**.**

Biometric Performance. *Design 2* relies on the error-correcting distance of the underlying fuzzy extractor to resolve biometric identity. The privacy features of the fuzzy extractor also hide match scores. It is more difficult to tune biometric error rates using traditional approaches, such as ROC and DET curves. This is because, in traditional biometrics, a floating-point value is generated as the score, which can then be evaluated against different thresholds on historic performance data to obtain an ROC curve. The ROC curve is not widely used in fuzzy extractor-based approaches. This is due to a number of reasons. Fuzzy extractor-based approaches hide match scores. Also, error-correcting codes used in FE transform the continuous score comparison to a discrete decode or receive no meaningful results. As such, fine-grained performance variations are directly observable with threshold changes in traditional systems. We can set and quantify “identity” decision differences as a function of the increasing difference in the distance metric between the issuance and enrollment images. A privacy biometric scheme must hide this information, because visible differentials in distance between probe and gallery images lead to vulnerabilities in terms of hill-climbing attacks (Adler, 2004).

Further work must be conducted in relation to fuzzy extractors for ROC curve analysis to allow for independent match thresholds between issuer and verifier. We assume that the system undergoes a configuration period in which biometric match success and biometric key gen are tuned to work together. A simpler approach for key

regen could involve placing the image binary and a random value as exponents in a Pederson commitment or as inputs to a cryptographic hash function and allowing the resulting value to be the regenerated key.

Interoperability. *Design 2* provides interoperability via image formats. Facial images are transferred between the user agent and the service providers. This approach to interoperability assumes the honesty of the verifier. A corrupted service provider can simply save transmitted biometric images. It is possible to place the entire fuzzy extractor key gen/regen mechanism on the user device so that the key gen and regen become the responsibility of the user agent. This approach is vulnerable when assumptions of user agent security are relaxed.

Design 2 requires agreement across the community regarding the error correcting codes of the fuzzy extractor. This feature of $Design\ 2$ can pose a problem in that it becomes impossible for verifiers to decouple their threshold choices from the wider community.

Adaptability. As discussed above, *Design 2* requires agreement across the community regarding the fuzzy extractor parameters, which naturally include the error-correcting distance. This feature of *Design 2* can pose a problem in that it becomes impossible for verifiers to decouple their threshold choices from the wider community. To appreciate the value of this adaptability, it helps to consider the need for an operational verifier to change match thresholds for illumination variances in the field or to accommodate different levels of image quality present in e-passport issuance.

5.5 Chapter Conclusion

This chapter described *Design 2* which was built on Brands digital credentials, including a fuzzy extractor-based RBR. This design delivers privacy properties not afforded by *Design 1*.

Based on Brands' digital credentials, *Design 2* provides selective show, credential composition, and a path towards untraceable transactions. The design, which is based on fuzzy extractors, has benefits of not requiring the issuer or verifier to store traditional biometric images. In terms of security analysis, the RBR is hidden in an extended Pedersen commitment, which appears to thwart known attacks against fuzzy extractors.

The fuzzy extractor design has some drawbacks with respect to biometric performance and interoperability. In terms of biometric performance, raw biometrics will always be the benchmark. Research on ROC curve analysis on fuzzy extractors is scarce compared with similar analyses of traditional biometrics. Performance reporting and ROC curve analysis in the area of privacy-biometrics are open areas of research.

Chapter 6 Collaborative Computing and Risk-balanced Cellular Access Control

6.1 Chapter Introduction

This Chapter departs from the themes of credential design in previous chapters and returns to the large-scale ecosystem perspectives introduced through EoS-UML in Chapter 3. We discuss the ability to run a distributed algorithm in a collaborative manner between stakeholder nodes, and propose one such algorithm - risk-balanced Cellular Access Control (rbCAC). The rbCAC algorithm is a subject screening algorithm that is distributed across the nodes in an ecosystem and allows the ecosystem to classify the subjects going through it as a function of the joint decisions of the distinct authorities. Within the rbCAC Cellular Access Control algorithm, we demonstrate how nodes may adapt their questioning based on perceived risk and route subjects to downstream nodes according to paths which reflect that perceived risk.

6.1.1 Scope

6.1.1.1 Contribution to the Collaborative Computing

1) **Proposed Environment.** Our execution model adds a multi-stakeholder risk/goal oriented perspective complementary to logical layer views of CPS architecture such as (Magureanu et al., 2013)(Coulter, et al. 2003), adding the layered

partitioning by generic processing element as a logical overlay, when helpful, to many of the physical architecture views of CPS (Vu et al., 2014)(Ma et al., 2017).

2) **rbCAC.** Definition of a novel risk-balanced Chained Access Control (rbCAC) algorithm. rbCAC is significantly different from the state of the art in risk aware access control in which work such as (Atlam et al., 2020)(McGraw, 2009)(Molloy et al., 2012). The rbCAC algorithm is multi-perimeter. Perimeters are cellular, recursive, additive continuous and dynamic. At this point in time, it seems the access control literature and the literature of CPS is largely focused on single perimeter models.
3) **rbCAC as a Collaborative Algorithm (CA).** In cyber-physical systems, recent work (Cao et al., 2020) looks at access control from the perspective of risk and access to facilities spatially and temporally separate. This is an advance in our direction: the chained nature and recursive nature of our zone model is not present in their work. We believe the application of EoS-UML to their application can yield benefit.
4) **Design by Smart Contracts (DbSC)** for ecosystem services allows the algorithm to express service levels expected variances, costs in terms of pre-conditions, invariants and postconditions. The introduction of confusion matrix-based output expectation with confidence allows statistical expectations to be set. AI techniques are involved in this currently (Jiang et al., 2008), however this has not been capitalized on by systems engineering. A natural thing to do, summarize the performance of a classifier using a confusion matrix – in the context of emergent behavior in CPS. As an innovation, we apply the confusion

matrix to design contracts for non-deterministic services. (Meyer, 1992) brought design contracts to the forefront in software construction. Design contracts have been considered within the field of SoS and CPS (Derler et al., 2013)(Battram et al., 2015) as a manner of reducing the uncertainty between independent components. Our contribution is to capitalize on statistical knowledge in the form of the confusion matrix and confidence interval in design contracts.

6.1.1.2 Exclusions

Self-optimization. This chapter does not implement network self-optimization. The Monte Carlo simulation is used to demonstrate processing. As well, game theoretic simulation and GRL integration are not conducted in this demonstration.

6.1.2 Setting Characteristics

Our setting features a collection of subjects in an ecosystem partitioned into regions in which multiple authorities control fixed-perimeter ROIs in a collaborative trust environment. We seek a design pattern which accommodates and capitalizes on:

1) **Measurable classification performance.** The overall classification of subjects into routes through the network is analogous to a machine learning problem. Subjects may have hostile or benign intention. The distribution of adversarial subjects is unknown. The authorities must make a routing decision (“permit”/”deny”/”further interrogate”/”escalate”) by subject’s estimating the risk posed by a subject of unknown intent based on observable attributes. The correlation of visible attributes to actual intent is not known. The algorithm should

offer the ability to measure and configure the performance of the subject-classification task in terms of correct and erroneous decisions.

2) **Objective-Driven, Mandate-Governed Stakeholders.** Within the ecosystem, stakeholders act in their own interests and according to their objectives. They provide data services according to the ecosystem regulations. Stakeholders will have different success criteria and risk tolerances. The proposed algorithm should include a policy-based approach which allows the stakeholders flexibility of processing according to their varied goals and risk thresholds.
3) **Progressive and Dynamic Decisions.** The ecosystem is composed of zones. A subject's progress through these zones naturally occurs in a gated sequential manner. This allows a staged decision to be made. The proposed algorithm should support progressive and collaborative decisions and risk categorizations processes between stakeholders.
4) **Cost and Uncertainty in Screening Sessions.** Subject risk is not uniform. Different screening techniques may be required in order to become confident of decisions. Depending on the risk, the data acquisition cost, and the information gained from unknown variables, the authorities may prolong an interview. Processing sessions and correct or incorrect decisions have associated costs and benefits. The algorithm should be able to balance cost, uncertainty, and risk.
5) **Multi-step, Collaborative, Adaptive Screening.** The authorities conduct screening on subjects sequentially. The decisions made about a subject are incremental, with the decisions of one authority affecting the subsequent processing of that subject by downstream authorities. A complete traversal of the

continuum consists of the individual path segments of an individual from checkpoint to checkpoint.

6) **Risk-Specific Processing**. The environment features dissimilarity of subjects, authorities, objectives and rewards. The distributed decision-making algorithm should permit risk-sensitive adaptive decision making, and this risk should be determined using authority-specific criteria. It should be possible to make risk-sensitive decisions based on transaction and context data.

6.1.3 Assumptions

1) **Credential Design.** A credential design offering stable security, privacy, and operational properties is assumed. Specifically, we assume the properties of unforgeability, tamper-resistance, and non-transferability of issued credentials (i.e., P1–P3) to be reliably present, without the risk of compromise. We assume data minimization and composability within the underlying credential scheme. This assumption allows the verifier and subject to ask and provide only those data elements required by the risk-scoring process. We assume data usability and classifier accuracy. We thus assume that the selected credential design provides authorities with data that they can properly analyze.
2) **Data Properties.** We assume that there are data patterns in subject attributes X that allow checkpoints to distinguish between honest and malicious subjects with an advantage greater than tossing a coin. We assume some subset X_j of a subject data attribute X exists and can be used by authorities to correctly guess

the given hidden intent Y_j, with a non-negligeable advantage over random guessing. If classifiers are used, we also assume sufficient training instances of historical observations such that classifiers may be suitably trained. We assume that the information contribution of the attributes {x,….x} is known by authorities with respect to their assigned subject intent risks.

Furthermore, we assume that classifier error and thresholds are known and can be meaningfully and dynamically updated to become stringent with respect to the attributes inspected and effective decision thresholds.

3) **Services.** This chapter assumes stable ecosystem directory and transcript services. For directory services ($\mathcal{D}$)**,** we assume the existence of a registry that can be used to select registered service providers based on credential offering and data quality (level of assurance). The directory can be used, for example, to resolve public keys, consult assurance levels, and determine possible costs of various attestations. For transcript services ($\mathcal{T}$)**,** we assume the existence of an ecosystem-public transcript on which public and leaked data can be viewed by transactors and spectators. We assume that this transcript information is reliable. We assume that $\mathcal{T}$ offers perfect security and privacy properties as specified by the designer of candidate collaborative algorithms.

4) **Points of Presence.** We assume subjects and local authorities have cyber-physical equipment (e.g., smartphone, cloud issuance, and kiosk) which they use within the ecosystem to represent their interests.

5) **Existence of a Planner.** We assume the existence of an overarching authority that establishes the ecosystem rules and policies and prioritizes trade-offs regarding optimal correctness, cost, and performance.

6.2 Building Blocks

6.2.1 Mealy State Machines

A Mealy machine is a 6-tuple (S, s_0. , Σ, λ, δ, ω), where S is a set of states, s_0 a start state, Σ the input alphabet, λ the output alphabet, δ a transition function $\delta: S\ X\ \Sigma \rightarrow S$, and ω an output function $\omega: S\ X\ \Sigma \rightarrow S$. For probabilistic FSM, the transition function is expressed as a probability $\delta: S\ X\ \Sigma \rightarrow$ S X [0,1], associating the transitions from a state with a probability. In Mealy state machines, ω is a function of the current state and the input. In contrast, the Moore model's output function depends only on the current state. We use the probabilistic finite state machine as a metaphor for modeling the concourse used in rbCAC . The algorithms $next\_state()$ and $update\_costs()$ correspond to transition function δ and output function ω of Mealy state models.

6.2.2 Quantitative Models

Our simulation approach is inspired by the discrete-time Markov chain. A Markov chain is a stochastic model describing a sequence of random variables (i.e., X_1, X_2, X_3, ...) with the Markov property, namely that the probability of moving from the current state to the next depends only on the present state, independent of previous states: $Pr(X_{n+1} = x \mid X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n) = Pr(X_{n+1} = x \mid X_n = x_n)$. In our

formulation, the subject carries its attribute data and hidden intent. The system adaptively uncovering the distribution and its optimal configuration is the goal. With the Discrete-Time Markov Chain (DTMC) approach, the subject does not carry data; instead, the distribution is assumed and reflected in the transition probabilities. We use this in Section 5.4 to demonstrate that the sensitivity of concourse output measures to changes in decision criterion can affect local and global performance. Our implementation presents a Monte Carlo simulation which uses a probabilistic mealy model implementation. By demonstrating how a risk sensitive orchestration harness can be made for discrete-time cost and risk calculations and goal assessments, and fitness function evaluation. This provides a basis a number of optimization strategies notably hill-climbing, simulated annealing and genetic algorithm, setting the stage for future work on self-optimization and self-adaptability. Future work may also examine various aspects of the Markov decision process and its suitability for modeling the rbCAC model.

6.3 Proposal

6.3.1 Concourse Execution Model

Figure 19 presents the concourse execution model. The subjects and points reside on the execution plane. Both are active, threaded, executable objects exhibiting non-deterministic behavior and behavior variance within the collaborative model effective in the trust setting.

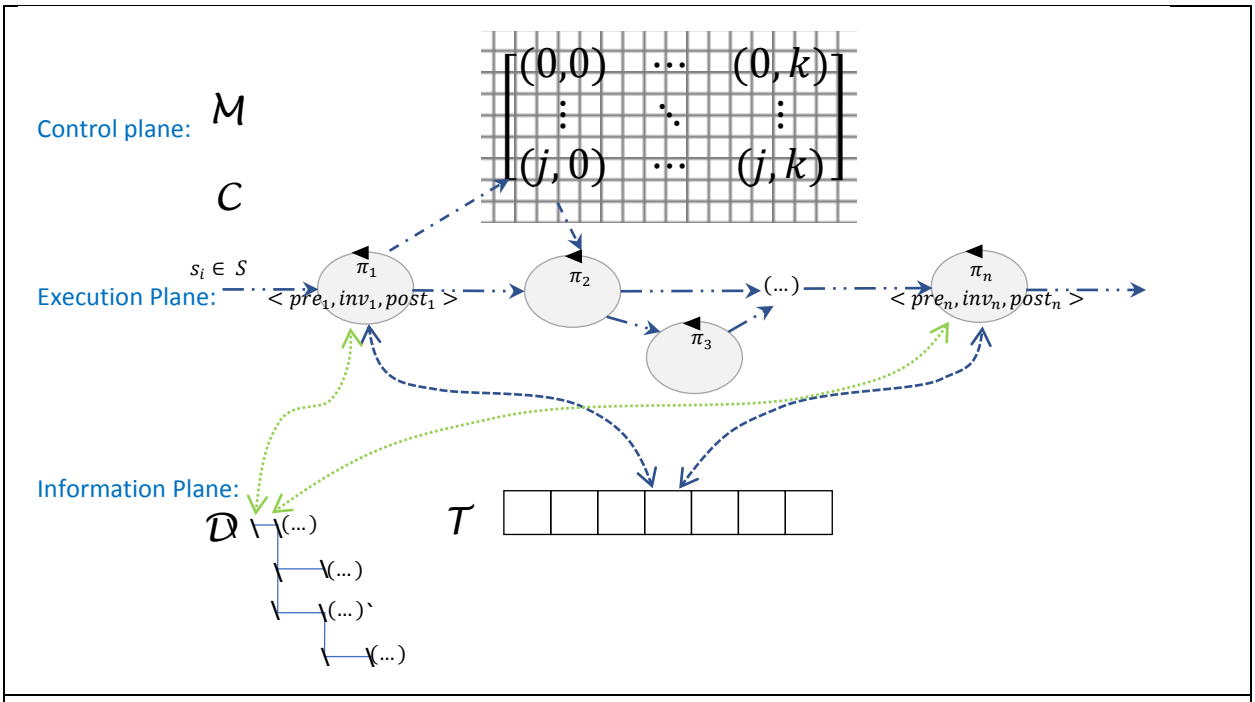

Figure 19) Cyber-Physical Ecosystem as a Distributed Processing Environment

Points of presence offer strong pre-conditions, invariants, and post-conditions to the environment. Since the control matrix decouples flow from nodes, the overall concourse and the individual trajectories are dynamic. Auxiliary objects $\mathcal{D}$ and objects $\mathcal{T}$ refer to the community yellow pages and whiteboard, respectively. The integrity of auxiliary objects may be maintained using distributed ledger technology.

6.3.2 Data Structures

The elements in our construction are shown in Table 8, below.

Table 8) rbCAC Concepts and Data Structure

Subject	Attributes vector X, intent vector Y
Concourse	G graph, Nodes n_{ij}, Adjacency list
Risk-Interval Map	T_{ij} transition definition from node i to node j
Lexicon	L lexicon attributes, relative attribute contribution, attribute acquisition cost
Directory	D Attestation provider's directory.

6.3.2.1 Subject

The subject is represented by two vectors: X attributes and Y (undisclosed) intention codes. As in Brands' Digital Credentials, vector X is an attribute vector and Credential C_{AI} granted by Issuer I to Alice was issued on those X. From the perspective of classification and decision making, the distributions of the X variables may be known, but those of the Y intent variables are unknown to the authorities. From an ML perspective, the ecosystem aims to learn a set of thresholds for which the flow of subjects through the network maximizes the objective function or best classifies for y' versus ground truth y.

A subject's traversal through the concourse becomes a progressive learning exercise in which the challenging team attempts to estimate hidden intent variables through the judicious application of risk and uncertainty sensitive screening. Publicly disclosed x attributes are subject to falsification by a deceptive attacker. The use of cryptographic credentials in which a trusted issuer certifies veracity of attributes helps mitigate the risk of false attribute claims from a malicious actor.

6.3.2.2 Objectives

The aggregate set of objectives G may contain objectives g_x and g_y, which are conflicting. However, it is the objective of the overarching authority a_0 to achieve a set of ecosystem risk screening questionnaires and thresholds which balances the objectives to maximize G, respecting the objective constraints g_i.

6.3.2.3 Concourse

The concourse is represented similarly to an adjacency list representation of a digraph. The concourse is a set of associations between each node and in the concourse graph and the corresponding risk interval adjacency list which maps from risk level to the label for the next state. Let concourse C be a dictionary that contains an entry for each state associated with an interval map, which lists the possible transactions from that state under particular risk levels.

6.3.2.4 Risk Classification Subject to Error

In general, all stakeholders conduct a risk assessment prior to entering a transaction. This risk assessment is subject to Type 1 and Type 2 errors. The expected value of this classifier performance, in terms of a confusion matrix and a confidence interval can be measured over time and posted as a service post-condition. Figure 20 shows a sample risk function with post-conditions on the return value as performance expectations for consumers of the service.

Figure 20 shows *R(…)* the method body for risk assessment with published expected performance. The function implements risk assessment in terms of a notional classifier. The focus here, is not on the specific risk assessment algorithm, but rather on

that that the non-deterministic function subject to classifier error is published to consumers with estimated performance measures.

Function R(subject, context):

y' = $\text{predicted}_{\text{cat}}$= assess_risk_category (subject, context)

return (y')

Post-condition: [X→Y PR(y' = Y[i])<= ε]

.

Figure 20) Risk Function with Classification Error

Risk is modeled as a classifier that predicts a risk class *y*' given subject data s, with error ε.

Referring to Figure 2-0, "post-condition: [X→Y PR(y' = Y[i])<= ε]" gives a market-place established "quality" to the estimated produced by the service. The veracity of the prediction can be expressed as $b' = (y' == Y[i])$, where Y[i] is the ground truth and y' is the prediction probability bounded by an error term ε. Expected Classification Error ε can be expressed in terms of the TPR, FPR and FPR, and FNR of a confusion matrix and the applicable confidence intervals. The post-condition acts as a service level indicator, and a buyer-beware notification to the marketplace. This measure expresses an expectation on provider's data analytics within a confidence band. This measure is a post-condition on the service. This post-condition may further be supported through reputation metrics provided semi-formally by clients and/or peers, as per the ALTM.

community's confidence on the veracity of the data contained in an attestation, represents a precise "Level of Assurance" or an "Information Assurance Level" (NIST 800-63-3) measure.

6.3.2.5 Risk Interval Map

The possible transitions from a state are represented in an interval-dictionary and defined as follows. Let partition interval $I = \{ t_0, t_1 \ldots t_n \}$ consist of a strictly increasing sequence of thresholds $t_0, t_1 \ldots t_n$ on the real number line with $0 = t_0 < t_1 < t_2 \ldots < t_n = 1$. A pair of the form $[t_i, t_{i+1}]$ is a subinterval of the partition I. For each subinterval in I, a next state label $n_i \in V$ is associated. The next state labels are stored in array L of size $n - 1$. The resulting data structure RM = (I, L) is a risk interval map which can be used to map from a quantified risk to the next state.

6.3.3 Algorithms

A discussion of algorithms $traverse()$, $policy\_stabilize\_risk(\ldots)$, $next\_state()$, and $update\ costs()$ follows.

6.3.3.1 Routing Through the Concourse: #traverse(…)

Algorithm $traverse()$ routes a subject through the concourse and returns the sequence of states visited, from start node to accepting terminal node.

```
traverse(subject, start_state, stt, C, pq_iv):
[1]  quit = false; node_name = start_state; tc traversal = []
[2]  while ( ! quit)
         risk_tiers = stt[node_name]
[3]      total_cost = C[node_name]
[4]      if (! size(risk_tiers) == 0 )
             (risk, uncertainty, interview_cost)
[5]              = policy_stabilize_risk(subject, risk_tiers,
     pq_iv)
[6]          node_name = next_state(risk_tiers[risk])
[7]          total_cost += interview_cost
         else
[8]          quit = true
             traversal.append([node_name, total_cost ])
```

```
return traversal
```

Algorithm 1. Concourse>>traverse

6.3.3.2 Authority Specific Risk-Interviewing: policy_stabilize_risk(…)

The algorithm to balance uncertainty, cost, and risk depends on the policies of the specific authority. Figure 21 presents an approach to balancing interview cost, uncertainty, and risk.

```
policy_stabilize_risk(s, risk_tiers, iv)
[1]   r = 1; u = 1; tc = 0
[2]   c1_iv_ok = 1; c2_budget_ok = 1; c3_risk_ok = 0;
      c4_uncertainty_ok = 0
[3]   nb = this.get_budget()
[4]   ut = this.get_uncertainty_threshold()
[5]   continue = 1
[6]   while (continue)
[7]         if(is_empty(iv))
[8]             c1_iv_ok = 0
[9]         else
[10]            a = peek(iv)
[11]            c2_budget_ok  = (nb <  tc + cost(a))
[12]        if (continue)
[13]                x = acquire(s,a)
[14]                X.append(x)
[15]                (r,u) = assess_risk(X)
[16]                c3_risk_ok = risk_tiers.includes(r )
[17]               c4_uncertainty_ok  = u > ut
[18]                tc += cost(a)
                    pop(iv)
[19]        continue = (c1_iv_ok  && c2_budget_ok  &&
```

```
                                ! (c3_risk_ok  |  c4_uncertainty_ok ) )
        return (r, u, tc)
```

Figure 21) Authority-specific Method: Policy_ Stabilize_Risk()

Algorithm $policy\_stabilize\_risk\ (s, tisk\_tiers, iv)$ receives as argument the subject, the *risk_tiers* definition and the information value ordering iv. The method engages in an adaptive interview-assessment loop (lines [4]–[18]) that interactively queries the subject (line [13]) until the incremental information gain of the next question versus the authorities $budget$, the perceived risk of the transaction relative to stakeholder risk tiers, and level of prediction uncertainty with respect to a transaction-specific confidence threshold ut (lines [10], [14], [15], [16] respectively) are satisfied. The cost for the interview tc is recorded (line [16]) and returned to the calling method, along with a risk classification r and uncertainty measure u. The cost of the next attribute is determined based on Directory contents on the domain lexicon. The loop continues as a function of the balancing conditions. These will change depending on the circumstances. In line[18] demonstrates the continuance criterion to be a Boolean predicate on the node's status and budget, and the perceived transaction risk and data uncertainty.

6.3.3.3 Transitioning Under Risk: next_state(…)

The algorithm names the state to which the traversal should be transitioned, given the current state and the perceived risk. Specifically, $next\_state(\ \ )$ in Figure 22 returns ℓ_{ir}, the name of the next state in the concourse to which traversal routed. The algorithm implements the Mealy model δ (see Section 6.2.1). The label of next state is returned

using an index lookup based on risk range. Specifically, the implementation uses a risk interval map to translate from risk interval to discrete risk index. First, concourse C is indexed using the current state s, which returns the interval map I dictionary of risk-sensitive transitions from the current state s. The return value ℓ_{ir} is then obtained by indexing into LI with index ir, the index of the subinterval within I which bounds the perceived risk r.

next_state(s, r):

[1] I = C[s]

[2] ir = min(i) (t_i <= r <= t_{i+1}) for risk r, and i ∈[0, ... , n]

[3] ℓ_{ir} = LT[ir]

[4] return ℓ_{ir}

Figure 22) Concourse method: next_state()

6.4 Demonstration on Air Traffic and Border Security

This section demonstrates a Monte Carlo simulation of rbCAC demonstrated on three ATBS scenarios.

6.4.1 Introduction

This experiment demonstrates how the risk thresholds and cost values of a concourse can be configured, how the performance of a concourse can be quantified, and its parameters refined. The model is a simplified version of the algorithm presented in the previous section. Rather than a risk assessment, checkpoints use probabilities to determine a subject's path. The distribution of risk and its classification are reflected in

the probabilities. Similarly, costing is simplified. A magnitude value representing cost/benefit to the system of a subject at a checkpoint or a materialized risk. These values are for illustration purposes only. Probabilistic thresholds are meant to be tuned by the ecosystem designer, or by the optimization harness (ex.: Genetic Algorithm or Simulated Annealing). Costs are meant to be discussed and refined with domain experts.

6.4.2 Visualizing Scenario Configurations

We examine three scenarios (Sc1, Sc2, Sc3) representing different levels of automation within ATBS to illustrate the difference that model parameters may have on total ecosystem performance. The scenarios are executed in a Monte Carlo simulation each on n = 1000 traveller-subjects.

6.4.2.1 Sample Configurations

Visualizations for three configurations are presented in Figure 23, Figure 24, and Figure 25. Sample cost/utility values and risk data for each scenario are provided in Table 9 and Table 10. Simulation results are provided in Table 12. It should be noted that the configuration and flow diagrams are both presented in EoS-UML.

Scenario 1: Manual issuance and verification. In the first scenario, the visa office is assumed to implement a manual Travel Authorization issuance process which is more costly, but perhaps more precise.

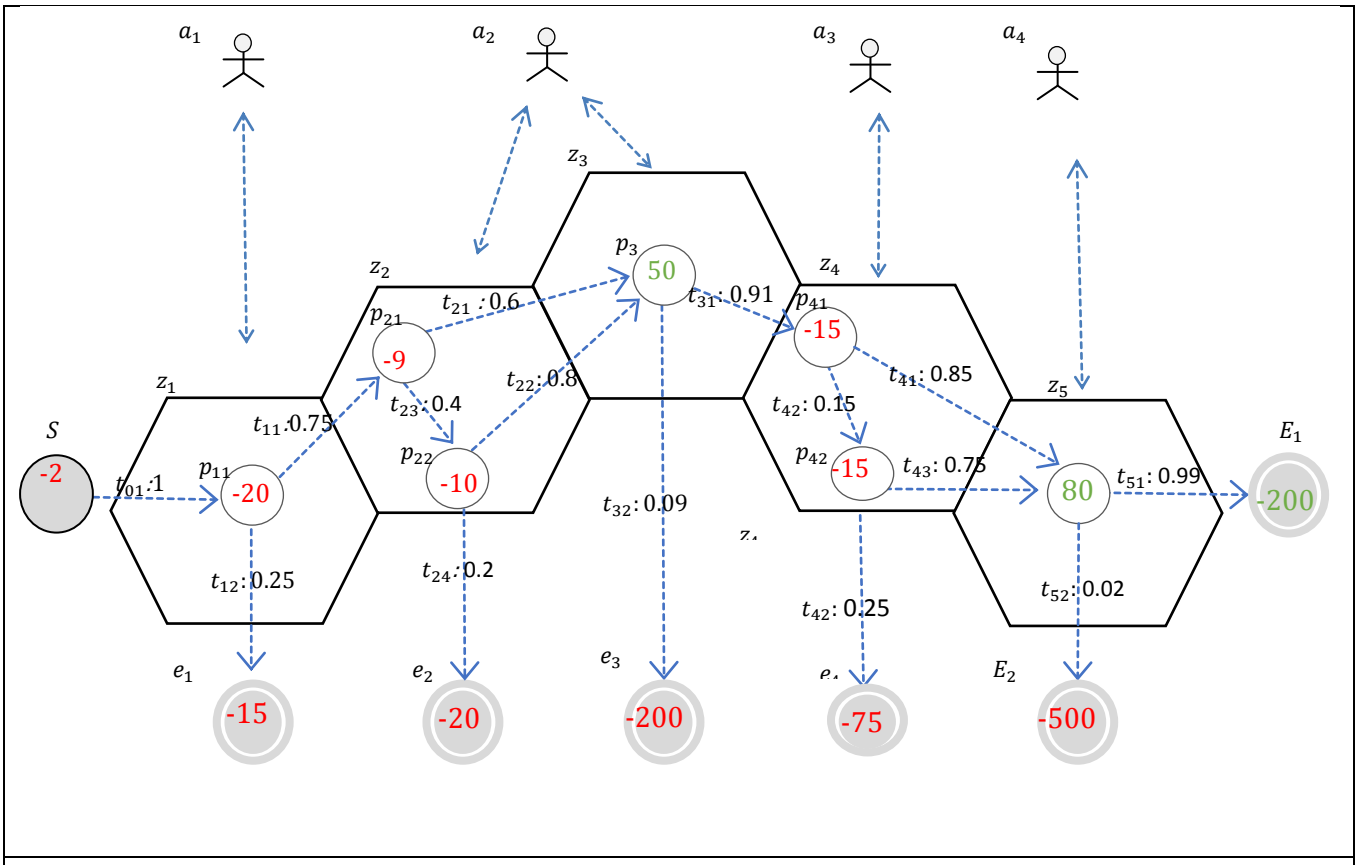

Figure 23) Scenario 1 Configuration Visualization

Similarly, the airline and border services are assumed to use traditional manual screening processes.

A new authority is introduced a_4, law enforcement of the country of destination operating in a new zone z_5 the country of destination. The addition of this zone and authority allows the scenario to show the complete picture of traveller flows, culminating in E_1 or E_2 in which an honest or malicious intent is revealed. Scenario 1 can be used to arrive at an estimated baseline for traditional processes.

To specify a configuration in Eos UML, transaction costs are placed on the nodes and probabilities or risks on the edges. Looking at Figure 23, for example, node S has cost of -2 which is incurred by each subject. Transition t01 from S to checkpoint p_{11} occurs with probability 1. checkpoint p_{11} has a processing cost of -20. At checkpoint p_{11}, 25% of the traffic is routed onto transition t_{12} where it ends up in defined exit e_1 incurring a penalty of -15. Seventy five percent of travelers are granted an mTA at p_{11} and proceed along transition t_{11} a preboarding check-in at p_{11}. The rest of the configuration can be interpreted in a similar manner.

Scenario 2: Automated issuance and partially automated verification. In the second scenario, the visa office uses a web-delivery system to assess mTA applications and issue them. The verifiers only implement partial automation of verification, with the airline using manual triage processes, and the border using automated and adjudicated processes.

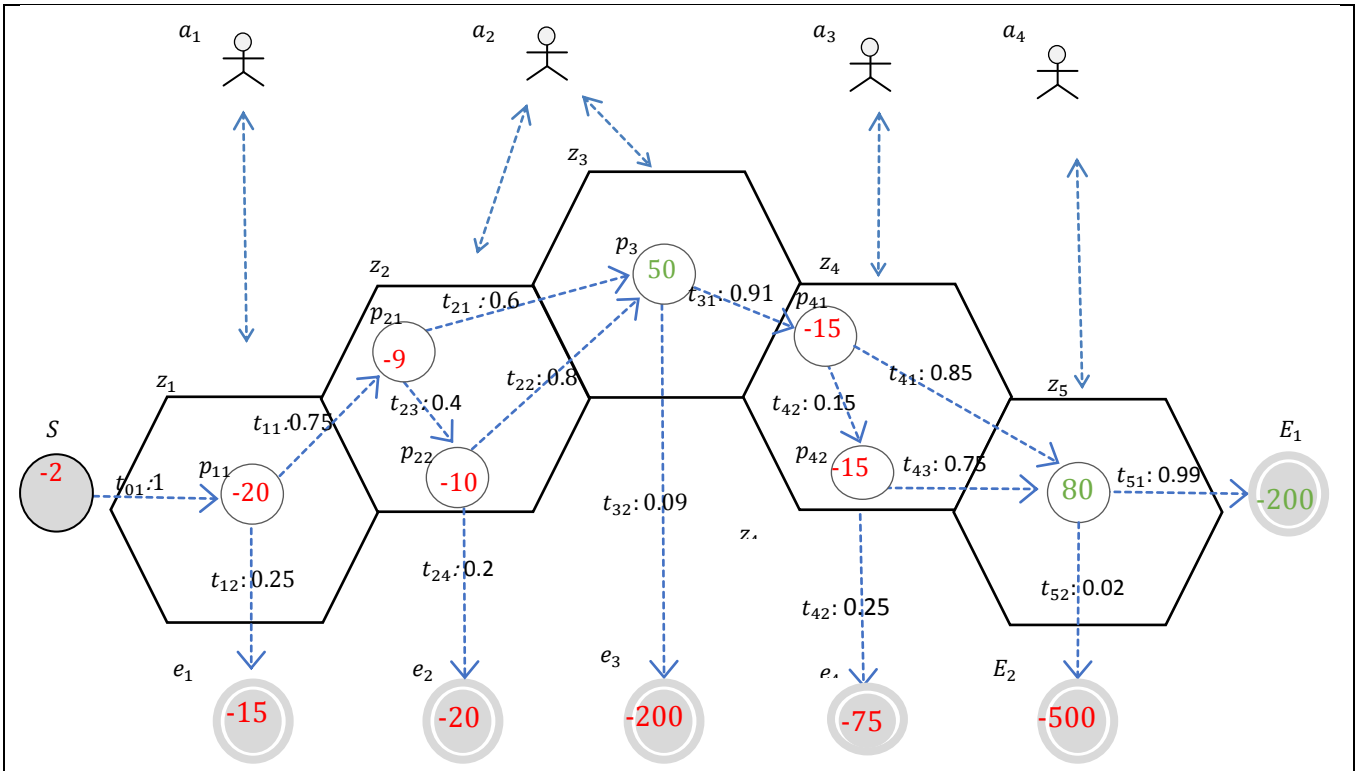

Figure 24) Scenario 2 Configuration Visualization

This scenario resulting in lower average checkpoint cost and improved risk-appropriate subject routing. Pre-board processes are still manual. There is some error due to manual data entry and the transaction time is generally slower.

Scenario 3: Full automation for issuance and verification. In scenario 3 issuance and both verification checkpoints implement automation. Verification functions include possible manual adjudication for exceptional situations. This fully automated scenario brings most benefits in efficiency and risk-appropriate routing.

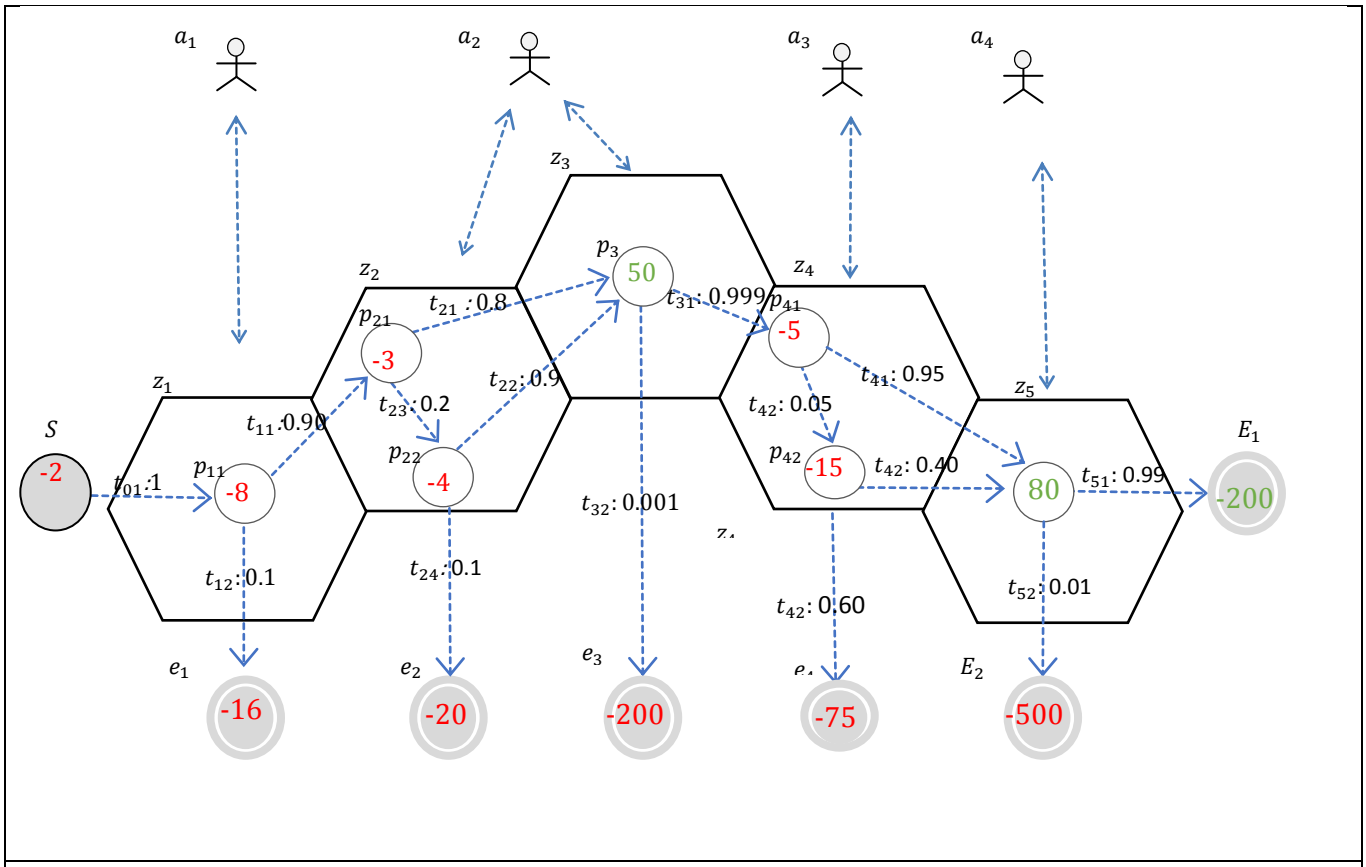

Figure 25) Scenario 3 Configuration Visualization

6.4.3 Estimated Costs and Probabilities

6.4.3.1 Comparative Costs for Scenarios 1 to 3

Table 9 shows the entry costs for each state under the three comparative scenarios. State S shows the same cost across all scenarios. For the mTA issuance step p_{11}, TA processing is more costly in *Sc1* due to manual processing.

Table 9) Comparative Cost Summary

Id	Name	*Sc1*	*Sc2*	*Sc3*
S	"Start"	-2	-2	-2
p_{11}	"TA issue"	-20	-8	-8
e_1	"TA denied"	-15	-16	-16
p_{21}	"preboard main"	-9	-10	-3
p_{22}	"preboard exceptions"	-10	-11	-4
e_2	"refused boarding"	-20	-20	-20
p_{31}	"in flight"	50	50	50
e_3	"in flight disrupt"	-200	-200	-200
p_{41}	"arrival primary"	-15	-5	-5
p_{42}	"arrival secondary"	-15	-15	-15
e_4	"arrival deported"	-75	-75	-75
p_5	"admitted"	80	80	80
E_1	"compliant stay"	200	200	200
E_2	"disruptive stay"	-500	-500	-500

automation, as is seen in the *Sc2* and *Sc3* entry costs for p_{11}. The costs for e_1 show that the disincentive to refusing a travel authorization marginally higher in the automated scenario as it becomes a service level target of the automation programme. At the pre-boarding checkpoint, p_{21} in *Sc1* and *Sc2* reflect manual processes. The entry cost of p_{21} increases slightly in *Sc2* because of the extra verification requirements mTA which in *Sc2* must be done manually. Cost drops significantly in *Sc3* due to automation. The entry cost of p_{22} increases slightly in *Sc2* for reasons similar to p_{21}. Costs decrease in Sc3 because automation reduces the number of passengers needing help. The cost of a refused boarding transaction stays the same in e_2 for all three scenarios. Similarly, the benefit of successfully boarding p_{31} is constant across all scenarios. The cost of an in-flight disruption at e_3 is constant across all scenarios. The cost of primary processing p_{41} decreases in *Sc2* and *Sc3* due to automation. The cost of secondary screening p_{42} and deportation e_4 . The immediate benefit of a successfully admission p_5 is constant across all scenarios. The benefit to the system of an honest intent traveler E_1 is constant across all scenarios, as is the cost to the system of a malicious traveler E_2.

6.4.3.2 Comparative Transition Probabilities

Table 10 shows the probabilities of transition from under each scenario. A mapping is assumed between risk and probability. All scenarios have the same starting probability at t_{01}. In t_{11}Sc1 experiences more rejections due to exactitude of manual processes. automated processes achieve target performance in *Sc2* and *Sc3*. For transition t_{12} rejection figures are consistent with the acceptance figures in t_{11} following a similar rationale. In t_{21}, only *Sc3* has automated processing at pre-board.*Sc1* and *Sc2* thus have lower volume transitions.

Table 10) Comparative Transition Summary

edge	endpoints	*Sc1*	*Sc2*	*Sc3*
t_{01}	(S, p_{11})	1	1	1
t_{11}	(p_{11}, p_{21})	0.75	0.9	0.9
t_{12}	(p_{11}, e_1)	0.25	0.1	0.1
t_{21}	(p_{21}, p_{31})	0.6	0.6	0.8
t_{22}	(p_{21}, p_{22})	0.4	0.4	0.2
t_{23}	(p_{22}, p_{31})	0.8	0.8	0.9
t_{24}	(p_{22}, e_2)	0.2	0.2	0.1
t_{31}	(p_{31}, p_{41})	0.91	0.91	0.999
t_{32}	(p_{31}, e_3)	0.09	0.09	0.001
t_{41}	(p_{41}, p_5)	0.85	0.95	0.95
t_{42}	(p_{41}, p_{42})	0.15	0.05	0.05
t_{43}	(p_{42}, p_5)	0.75	0.4	0.4
t_{44}	(p_{42}, e_4)	0.25	0.6	0.6
t_{51}	(p_5, E_1)	.98	.99	.999
t_{52}	(p_5, E_2)	0.01	0.01	0.01

In t_{22}, a small percentage of travelers are forwarded to manual processing in an automated pre-board scenario. In t_{23}, *Sc2* manual processes and the new mTA processing requirements negatively affect productivity. Costs at *Sc2* increase, the number of persons permitted to board decrease due to efficiency and more stringent immigration requirements. For transition t_{24}, scenarios *Sc1* and *Sc3* experience some false negatives due to volume and its impact on efficiency. *Sc3* performs most smoothly because only exceptional cases are processed manually.

Transition t_{31}, which represents eventless passenger flights. In the example, it is inversely proportional to t_{32}, passenger disruptions in-flight. Transition t_{32} shows decreases in disruptions, especially in Sc3 which features automated e-passport

verifications prior to boarding. For t_{41}, manual processes in *Sc1* result in more referrals to secondary. Referrals to secondary in t_{42} decrease with automation. In t_{43}, the number of low-risk travelers that are referred to secondary decreases with pre-screening measures. The number of resultant referrals to secondary also increases in t_{44}, as a result of more precise risk assessment. In t_{51}, the number of honest travelers admitted to the country of destination shows an increase. The number of malicious travelers are admitted t_{52} commensurately decreases.

6.4.4 Visualizing the Flows

The rbCAC simulation can be conducted given probability and cost information. Figure 26 shows the traveler flow for Scenario 3 in simulation of n = 1000 travelers.

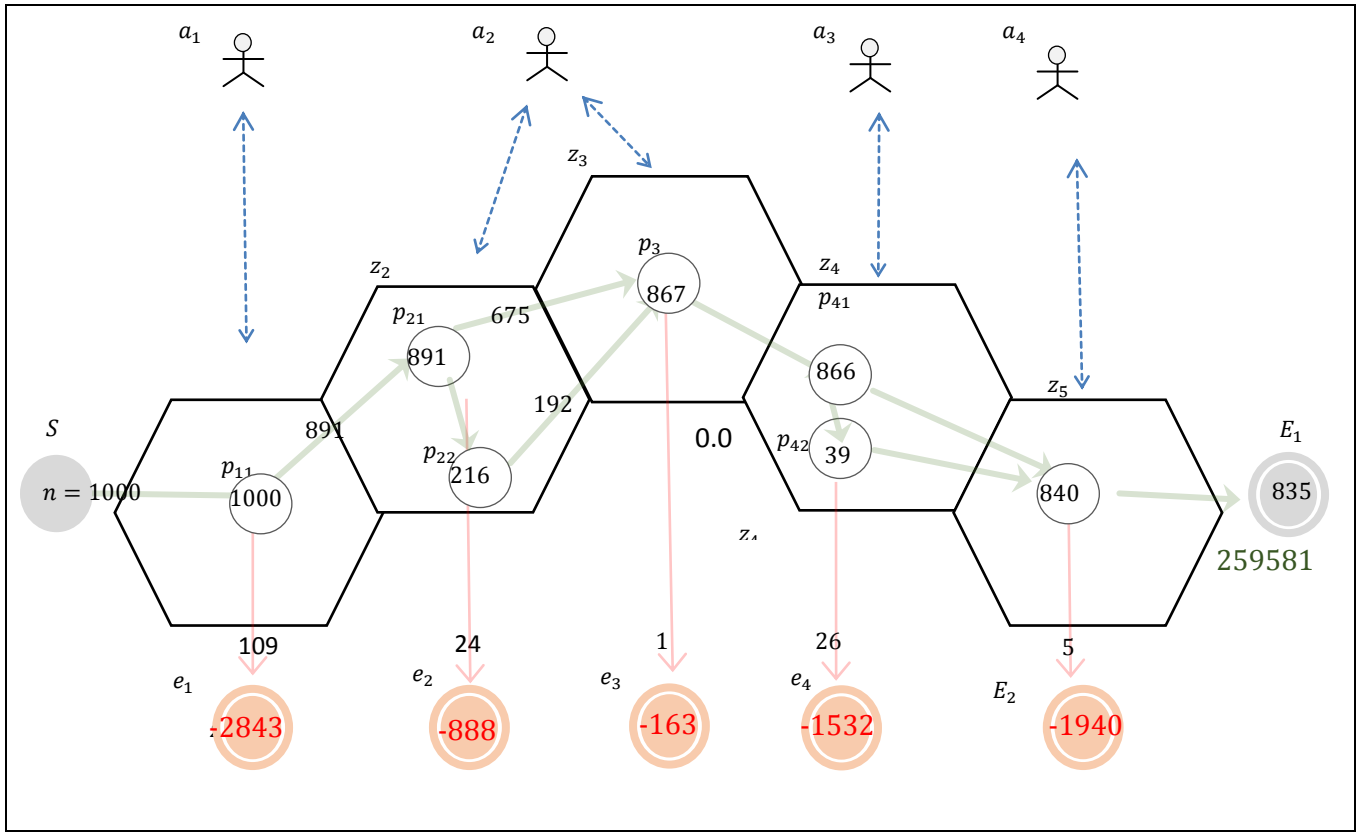

Figure 26) Scenario 3 Flow Visualization

The checkpoints in Figure 26 is made with the node visit data in Figure 27. From left-to-right, in Figure 26, 1000 subjects pass through the concourse, at progressive checkpoints some subjects are routed onwards, others are denied passage, or routed to further screening processes. This dynamic can be understood by examining in z_2, for example. Of the 1000 subjects entering the concourse, 109 were denied an mTA and 891 proceeded to pre-boarding.

Of those 891 passengers, 675 were permitted to board the plane without further screening however 216 were further interviewed to determine if they would be permitted to board. Of those 216 passengers, 192 were allowed to board but 24 were denied boarding.

```
'S_start': 1000,
'p11_TA issue': 1000,
'p21_preboard_main': 891,
'p31_inflight': 867,
'p41_arrival_primary': 866,
'p5_admitted': 840,
'E1_compliant_stay': 835,
'p22_preboard_overflow': 216,
'e1_TA denied': 109,
'p42_arrival secondary': 39,
'e4_arrival_deported': 26,
'e2_no_board': 24,
'E2_disruptive_stay': 5,
'e3_inflight_disrupt': 1
```

Figure 27) Checkpoint Frequencies for Scenario 3

This resulted in a total of 867 passengers in flight, one of which caused a disruption. Thus 866 passengers remaining, whose intent is not revealed. They are channeled to primary inspection, which in Scenario 3 uses an electronic screening transaction with a kiosk point of presence similar to discussed in Chapter 5. Of those passengers 827 are permitted entry into the country of destination while 39 are routed to secondary processing for further screening. Twenty-six of those passengers are not permitted into the Country of Destination and are detained/deported. Of those travellers permitted into the country of destination, 5 reveal during their stay a malicious intent, whereas 835 are well-intentioned tourists. Cost-Utility calculations and confusion matrices can be derived using these data and are discussed in the next section.

6.4.5 Metrics

This section provides assists in interpreting the results. The utility-cost calculations and confusion matrix information are the key fitness indicators to be used in determining the effectiveness of an rbCAC configuration as a classifier of a stream of subjects.

6.4.5.1 Sample Cost-Utility Calculations

Each subject's traversal through the concourse results in a net benefit or cost to the system. The performance of a checkpoint and the entire concourse can be expressed using a variety of metrics. Subject-flow and Cost-Utility are presented here.

Subject-flow. Subject-flow calculations use the traveller volumes to show how many subjects crossed a point in the concourse. Referring to Figure 26, these values are presented as positive magnitudes on nodes and optionally edges. Referring to z_5 in Figure 26, for example, we can infer that 840 travellers entered p_5 and 835 terminated in E_1 whereas 5 terminated in E_2. Flow information is cumulative. The 840 travellers that entered p_5 consist of $(866 - 39) = 827$ which entered p_5 from p_{41} and $39 - 26 = 13$ which entered p_5 from p_{42}.

Cost-Utility. The Cost-Utility measure can be used to characterize a subject's traversal, a checkpoints net performance, or the net performance of the entire network. In Figure 26 we see the total cost of -888 assigned to node e_2, "boarding denied". This amount can be calculated adding the traversal cost of each subject whose terminal node is e_2. The path to e_2 passes through nodes (S, p_{11}, p_{21}, p_{22}, e_2) having entry costs of $-2, -8, -3, -4$, and -20 respectively. Per the layout of the concourse, this is the only possible path from S to e_2. The path has a total cost

of -37. From Figure 26 we see that 24 subjects have traversals ending in e_2 the total cost assigned to e_2 is thus $-37 * 24 = -888$. Other costs can be calculated in a similar manner with the note that there may be more than one path to an end node. In those cases, the cost of each path should be added. Two paths, for example, lead to e_3: (S, p_{11}, p_{21}, p_{22}, $p_{3,}$ e_3) and (S, p_{11}, p_{21}, $p_{3,}$ e_3). Figure 26 indicates that only 1 traversal ended at e_3 with a traversal cost of -163. We can infer that in this case, the traversal followed the second path.

Confusion Matrices. Similar to the subject-flow and cost-utility measures discussed above, the confusion matrix can be calculated on an aggregate level for the concourse or at a local level, for each checkpoint. The data required to produce the confusion matrix for p_5 is included in Figure 26 with the additional assumption that 3 of the 26 people deported in e_4 were honest travellers, and thus false negatives.

Table 11) Confusion Matrix for Admissibility Decisions

Condition / Prediction	True	False	Total
True	835	5	840
False	3	23	26
Total	838	28	866

the standard classifier metrics can be calculated from the confusion matrix data in Table 11 Such confusion matrices can be produced for each node locally, and also for the entire concourse.

6.4.5.2 Comparative Flow and Utility

Table 12 shows the number of passengers channeled to each exit in all three scenarios, with the accompanying cost or utility at each exit.

Table 12) Comparative Flow and Utility

	Magnitude	e1_TA denied	e2_no_board	e3_inflight_disrupt	e4_arrival_deported	E1_compliant_stay	E2_disruptive_stay	Total
Sc1	Travelers	267	51	53	20	581	28	1000
	Utility	-9879	-3111	-9753	-1800	161819	-11838	125438
Sc2	Travelers	112	66	74	24	695	29	1000
	Utility	-2912	-3234	-12796	-1666	209945	-11506	177831
Sc3	Travelers	109	24	1	26	835	5	1000
	Utility	-2834	-888	-163	-1532	259581	-1940	252224

The “Traveler” values correspond to the number of subjects that exited at each named exit. The total number of travellers is 1000 in each scenario. The values in the “Utility” row correspond to the total cost or benefit produced by the traversals terminating at the given exit. Referring to scenario 3, for example, we see that 24 subjects exited the system at $e2\_no\_board$ for a total cost of -888. These values are calculated as described in §6.4.3.1.

Total utility for each configuration is given in the last column. Table 12 shows a progressive increase in utility from Scenario 1 through Scenario 3 as automation is added and progressive risk assessment is implemented. The highest utility of 252224 is demonstrated in Scenario 3 in which pre-board and custom utilize automated passport and facial recognition systems.

6.5 Conclusion to Chapter 6

This chapter presented an ecosystem algorithm for distributed access control. The algorithm is dynamic, altering the questionnaires of checkpoints based on the perceived risk of the subject of interest. The algorithm is also progressive, allowing the risk assessment decision of an authority to build on the information learned by other authorities. Checkpoints are connected in a directed graph, and the resultant structure is treated as a state machine. Subjects are processed by the state machine, one checkpoint at a time. Each checkpoint asks the questions it deems appropriate to make a routing decision for the subject at hand. A subject's passage through the network is thus a result of the progressive decisions made at each checkpoint. Checkpoints may grant credentials to subjects and/or write data to the system transcript that assists downstream checkpoints in making their routing decisions. Subjects have a hidden set of intent variables that determine their honesty (or malice). Both correct and incorrect decisions incur costs. The network's goals are to maximize net benefit while attempting to satisfy the local goals of each authority. This chapter presented the algorithms that are required of the state machine and local checkpoints. Risk and attribute costs are relative to the goals of the stakeholders and are thus ecosystem dependent. This chapter provides algorithms which assume attribute costs based on information gain.

Chapter 7 Selected Topics

This chapter discusses selected topics in the main areas of contribution: Ecosystems of Systems (EoS), their engineering and the Cyber-Physical Ecosystem Threat Model (CPE-TM): the comparative design of cryptographic credentials (CC) and their properties; and the ecosystem as a collaborative processing environment with emphasis on the risk-based Cellular Access Control (rbCAC) algorithm.

7.1 Topics on Ecosystems of Systems

Topics of discussion within EoS include the SoS Taxonomy and other possible types of ecosystems, EoS-UML, and the CPE-TM and its application.

7.1.1 Ecosystem Design by Smart Contracts

Design by Smart Contracts (DbSC) allows functionality in the composite-system to be rendered more robustly as a function of functionality presented by component-systems. Recall that EoS-UML uses pre-conditions, invariants, and post-conditions for PoP to enhance the confidence of that Stakeholders engaging in transactions. Remember that PoP are transaction endpoints, representing their stakeholders in value interactions. These interactions bind component-systems. DbSC allows this combination to occur in a scaffolded manner, imparting confidence to transaction participants. This idea builds on Design by Contract (DbC) introduced in the Eiffel Programming Language (Meyer, 1992). Design Contracts have been proposed in the field of CPS (Derler et al., 2013)(Battram et al., 2015). In general, the proposed contracts are crisp Boolean

functions, not able to not capture the non-deterministic nature of human in the system of STS and CPE. We propose placing these contracts in auditable persistence such as a distributed ledger. In the case where a service has statistical uncertainty, we propose that the DbSC contract terms express the statistical expectation and confidence of the outcome with respect to a confusion matrix. Recall from Chapter 3 that points of presence serve as the representative of stakeholders on the terrain. Stakeholders collaborate in an arms-length manner, exhibiting a healthy “méfiance” (somewhere on a spectrum between blind trust and counter-productive mistrust). The pre-conditions, invariants, and post-conditions aspects of EoS-UML can help ensure accountability and thus reduce friction to collaboration. In a smart design contract approach, concourse nodes are equipped with public sets of boolean predicates which outline, for example, the technical entry requirements, processing expectations, and exit assertions which are guaranteed by the PoP. An example of a post-condition on a risk scoring method as was presented in in Figure 21 on page 145. Pre-conditions, post-conditions and invariants in these contracts can assist marketplace dynamics forming the basis for comparative service pricing and purchase. Furthermore, these serve the basis for quality of service reporting, audit/dispute resolution, and the enforcement of compliance to obligations. For example, GDPR compliance and enforcement may be assisted by providing automated manners to electronically summarize, monitor and enforce service obligation workflows.

7.1.2 Other Types of Ecosystems

Cryptocurrency ecosystems (CCE) were, arguably, the first instance of CPE to become evident. The multiplicity of stakeholders is clear: cryptocurrency traders, block miners, and the observing public. The credentials are the verifiability of signatures and the traceability of sufficient funds. The medium of exchange is the bitcoin itself, based on the value assigned to it. The transcript is the blockchain of transactions in which a transaction provides <from, to, amount> information with <from> providing observers the ability to aggregate transaction balance to verify sufficiency of funds for a transfer. In terms of properties, security is fundamental, particularly tamper-resistance and the unforgeability of signatures. The cumulative design of the blockchain guarantees that the transaction history cannot be altered. The properties of the public key signature system guarantee the integrity of the transaction data. Privacy is limited. To calculate sufficiency of funds, transaction traceability to an account is required. This is achieved through the public key, which becomes a type of trading pseudonym.

It is when we look at operational properties that things become interesting. Fungibility is an interesting case to consider, namely the ease with which a cryptocurrency can be converted to a traditional fiat currency. The property of fungibility highlights a theme underlying this thesis, i.e., that new properties will be identified as ecosystems evolve, and that these properties are centered on meeting the requirements of their stakeholders. The relevance of target properties is a function of the joint requirements of the stakeholders.

Is a CCE cyber-physical? Arguably, it is. The properties of unforgeability are achieved through a miner's ability to solve a computationally intensive problem that

amounts to a brute force search for which no useful preprocessing tasks are known. A miner's advantage comes from having access to low-cost electricity to power a computer, which is a geophysical advantage. Furthermore, the property of fungibility to fiat currency also reflects a cyber-physical nature.

In many ways the cryptocurrency problem is a simpler one than that of social services. In Bitcoin, pseudonymity is satisfactory. Privacy concerns the cryptographic treatment of a few fields. Validity is a Boolean expression on the sum of a numeric field. With social services transactions and transactions with a more involved lifecycle graph, on the other hand, the issues of privacy and accountability become more important. This, in part, is why the present work has chosen to focus on social services.

7.1.3 Trust Models for Ecosystem Participants and Adversaries

This thesis uses an iterative Monte Carlo simulation to demonstrate numeric results. This allows a hill climbing path towards optimization and analysis of convergence. A game theoretic model is possible and should be interesting in analyzing stakeholder interaction and system dynamics.

The ATBS scenario we have focused on represents a collaborative ecosystem actors is a cooperative one, and in game theoretic terms might be best represented as a Stackelberg or Cournot equilibrium (Cournot, 1838)(Basar and Cruz, 1981)(Sherali et al., 1983). As the dynamics between the ecosystem participants becomes competitive, the dynamics are adversarial. In game theoretic terms, this competitive dynamic is illustrated by a Nash equilibrium (Cruz, 1975)(Basar and Cruz, 1981)(Myerson, 1978).

The Nash game theoretic model can be used to illustrate the dynamics between the entire ecosystem and the adversarial attacker. As honest participants and entire regions become subsumed by the CPE_TM Attacker, the overall dynamic moves from a collaborative one between honest stakeholders with a possible win-win end game, to adversarial model with a zero-sum end game. The modeling of ecosystems in game theoretic terms is identified as an area for future research.

7.1.1 Privacy Properties

Unlinkability. *Design 1* offers no blinding function, so the same credential is sent over the wire in multiple uses. Over HTTP, the multiple presentations of the credential are traceable because the credential remains identical. The HTTPS protocol encrypts traffic such that the credential cannot be seen by eavesdroppers. Regardless of transport encryption, the encrypted credential itself is identical over multiple presentations and is thus traceable by colluding recipients.

Design 2 provides single-show unlinkability at the cryptographic primitive level. In Brands' digital credential scheme, the blinding function transforms the credential-signature pair used in verification such that they cannot be correlated to the issuance transaction, however the credential can not be used more than once, otherwise it becomes traceable. Furthermore, in Brands' digital credential scheme, attributes are sent in cleartext to the verifier. As such, it may be possible for an attacker to correlate transactions based on the attributes sent. Conducting Brands' protocol over HTTPS

solves these problems. If multiple show support is required at the crypto primitive level, anonymous credentials can be used.

Biometric Privacy. *Designs 1 and 2* both offer biometric privacy, which relies on the honesty of the verifier while engaged in the transaction. This assumption is required because in both designs actual biometric images are sent to the verifier, and the expectation is that they are deleted after use. A corrupt verifier may ignore this convention ("No-Persistent-Biometrics") and save the biometric information. In an DbSC scenario (§7.1.1) such a verifier would be in conflict with the contract's invariant conditions.

Design 2 offers a solution to simply trusting the service provider to not save biometrics. In *Design 2* , fuzzy extractor regen() functionality occurs on the kiosk. If this responsibility is moved to the user device, stronger biometric privacy may be achieved, since key regeneration would occur on the device. For this to work, the user device must be trusted to be well-behaved in terms of data privacy.

Selective-Show. *Design 1* does not permit a selective show. All attributes are encrypted, signed, and transferred in one package produced by the issuer for the holder, stored, and then sent verbatim from holder to verifier. *Design 2*, in contrast, allows a selective show: the subject may divulge a subset of the certified attributes to the verifier, depending on the needs of the latter. The EbC approach of *Design 1* (like the DTC and the mDL mentioned in Chapter One) has no feature of data minimization. *Design 2* is at the opposite end of the spectrum. Not only are the data minimization characteristics of *Design 2* attractive to the subject from a privacy perspective, but they may also be attractive to the service providers in the ecosystem. From the service

provider perspective, data minimization and credential composition allow right-sizing the information gathered from the subject. This may help minimize the data a service provider holds and thus reduce data liability from a honeypot/data-exfiltration perspective. From the service provider's perspective, data minimization allows the data transfer and storage costs to be minimized to only, and exactly, the attributes required.

Composability. *Design 1* does not offer composability. If a verifier requires attributes contained in two separate credentials, both credentials must be sent in the verification protocol. Since *Design 1* does not feature selective disclosure, this implies all the attributes contained in both credentials must be sent to the verifier. After decryption, the verifier has access to all attributes in both credentials. *Design 2* allows for composability. In both Brands' and Camenisch and Lysyanskaya's AbCs, proofs can be constructed which combine attributes from different credentials. Like data minimization, as noted above, this feature is attractive from a privacy perspective. Furthermore, it allows flexibility: as new issuers and credentials are added to the ecosystem, their attributes can be combined in proofs, to the benefit of issuer, holder, and verifier.

7.1.2 The Execution of an Attack

The cyber-physical attack (CPA) can be specified as a parallelized script. To be effective, each step in an attack must be executed by properly qualified personnel, on appropriate equipment, and at the right place and time. The StuxNet attack offers an example: it required infiltration by foot personnel on a USB key, the appropriate

preliminary dissemination hosts of a Windows operating system, centrifuge equipment, and a TCP/IP connection to call home. Consider the CPE depicted in Figure 28.

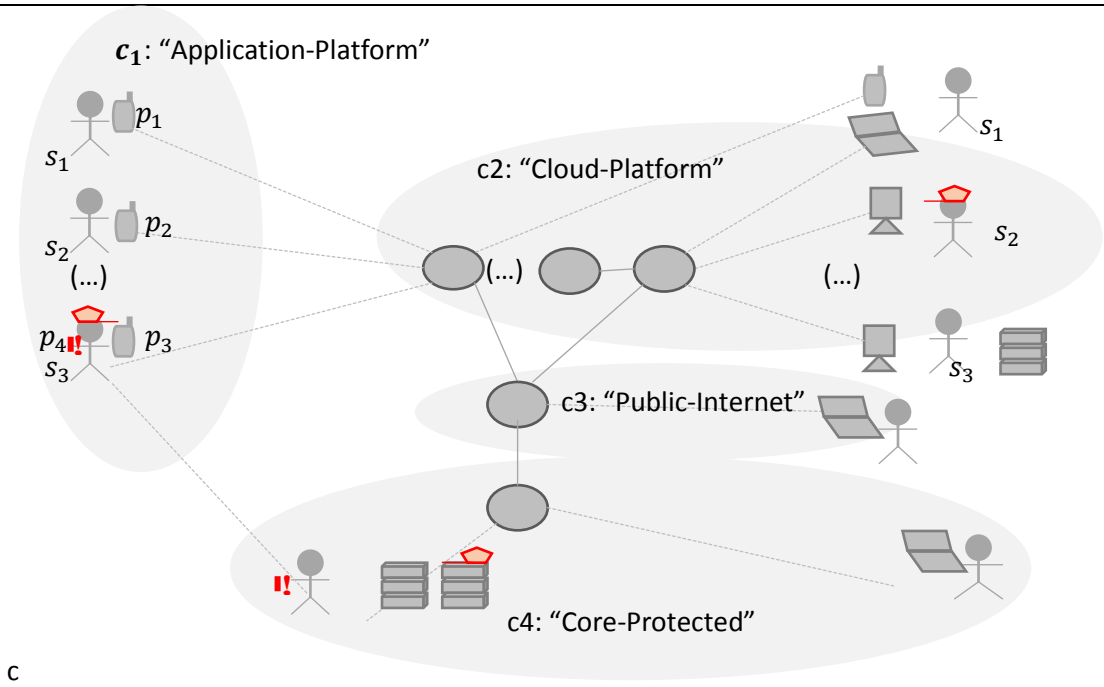

Figure 28) Cyber-Physical Context Tree

Let us represent the environment over time as a set of trees $\boldsymbol{T_t}$ over time $\boldsymbol{t} \in [1, \boldsymbol{s}]$, with s the number of environment samples taken. Now let $\boldsymbol{A_{t_0}} = \{\boldsymbol{s_2}, \boldsymbol{s_3}\}$ be the non-empty set initializing the terrain attacker. Assume the UML profile of Chapter 3 has been augmented with icons for smartphone, laptop, desktop, and server stacks as supporting systems and a red rectangle as a corrupted peripheral device. Consider the attack-script as a computer program – a FIFO queue in which instructions <$Context - Operator, Instruction, Payload$> are placed. Executing an attack-script consists of processing instructions by directing them to the appropriate $Context - Operator$ pair with appropriate forks and joins. Substituting stub/mock operators and contexts can allow for pre-launch testing of an attack.

7.1.3 Visualization of Attacks

An attack can be visualized by presenting the steps in a scripted attack and animating the attacker's progress in a frame-by-frame manner on an EoS-UML model. This approach can be used to visualize possible test attacks on a modelled target. Offensive attacks may be simulated by modelling the real-world SoS using stubs to represent RoI and protected resources and then animating an attack script. Figure 28 provides an example of a three-step attack which included creation of a corrupted token (ex. a USB Stick), its dissemination to an honest insider, and finally to successful penetration of c_4 through the infection of a protected server.

7.2 Properties of Comparative Credentials

This section contrasts *Design 1* and *Design 2* from Chapters 4 and 5 in terms of the security, privacy, and operational properties previously outlined.

7.2.1 Security Properties

Non-transferability. Both *Design 1* and *Design 2* provide a form of non-transferability using fuzzy extractors. The properties of **Unforgeability** and **Tamper Resistance** are delivered by the digital signature schemes of each design. The standard DSA scheme is proposed for *Design 1*, and Brands' digital signature scheme for digital credentials is proposed for *Design 2*.

7.2.2 Operational Properties

Data Reliability. Both *Design 1* and *Design 2* offer data reliability due to the security properties of non-transferability, unforgeability, and tamper-resistance of the credential and the use of actual attribute values or optimized representations in support service requests**.** In *Design 1,* the attribute values are transferred to the verifier. In the AbC approach, it is not necessarily the case that the attribute values are transferred. In *Design 2*, based on Brands' digital credentials, the values are also part of the service request. In an approach using Camenisch and Lysyanskaya's anonymous credentials, literal data values are, by default, not transferred. Rather, the truth of predicates in zero-knowledge proofs is what the prover communicates to the verifier. This has powerful potential if correctly used. If a more general approach is required, literal values may be sent as part of anonymous credential proofs as well.

Biometric Performance. *Design 1* offers a classifier-based approach for biometric matching whereas *Design 2* relies solely on the error-correcting distance of the underlying fuzzy extractor to resolve biometric identity. *Design 1* allows service providers to select a biometric matcher and tune biometric error rates using traditional approaches such as ROC as well as DET curves to achieve the operationally required performance. The use of fit-for-purpose classifiers in *Design 1* with the associated biometric distance metrics and performance measures has yielded more reliable performance than the error-correction approach of fuzzy extractors in our empirical studies. As described in §4.4.4.1, the use of both matcher and fuzzy extractor in *Design 1* can lead to problems if the match decisions of the biometric matcher and the fuzzy extractor do not agree.

In the field trials which accompanied *Design 1* , it was observed that commercial matchers perform more reliably than fuzzy extractors. The field trials were not conclusive; however, the results suggest that further work is required in fuzzy extractors for ROC curve analysis and to allow configurability and independent match thresholds between issuer and verifier.

One potential fix for this problem is the implementation of a key-release pattern. On issuance, the biometric key is a hash, or a Pederson commitment of the passport image combined with a random value. On verification, if the classifier confirms a biometric match, the key is regenerated using the image binary and the supplied random value.

Interoperability. Both *Design 1* and *Design 2* provide interoperability from the point of view of biometric matching. Both designs transfer facial images over the air between the user agent and verifier checkpoints. For the purposes of matcher comparison and fuzzy extractor keygen, the images are transferred to the verifier, who must be trusted to dispose of them after comparison. This level of interoperability thus comes at a certain price in terms of biometric privacy, as an assumption of verifier honesty is required.

It is possible to place the entire fuzzy extractor keygen/regen mechanism on the user device. Under this approach, image transfer to the verifier is no longer required; however, it results in the biometric performance issues described above. If we momentarily ignore biometric performance issues, the interoperability requirements of fuzzy extractors become obvious. *Design 2* requires agreement across the community regarding fuzzy extractor parameters, which naturally includes the error-correcting distance. This feature of *Design 2* can be problematic in that it becomes impossible for

verifiers to decouple their threshold choices from the larger community. This approach may lead to an interoperability and flexibility problem because all verifiers must be satisfied with the same error-correction parameters.

As discussed in §4.4.4.2, in its basic form, *Design 1* suffers from another issue that may perhaps be categorized as “interoperability”. This issue pertains to envelope encryption and decryption. Since, by design goals, for ease-of-adoption *Design 1* uses commodity encryption, the scheme uses point-to-point encryption/decryption between two parties. This means that in the verification pipeline – where the airline must first verify and then the border must check again – both authorities require the private decryption key. A simple solution would be to distribute the key to both verifiers; however, this may go against standard policies or best practices. Other potential solutions include broadcast encryption or attribute-/policy-based encryption. Using these, however, goes against the desire to use commodity cryptography exclusively.

Design 2, based on Brands’ digital credentials, offers strengths but also weaknesses in this “two-verifier” scenario. The strengths and weaknesses lie in the extended DL-rep commitment, proof syntax, and blinding function of Brands’ credentials. First, the extended DL-rep commitment encases the attributes in such a way that they are hidden from curious eyes and thus very strongly encrypted. Second, the proof syntax of Brands’ digital credentials sends attributes in cleartext across the wire. This is a drawback, as it is a source of possible transaction linking and de-identification. Finally, the signature scheme and blinding function of Brands’ digital credentials only permit a “single-show” of a credential. Thus, in the “two-verifier” scenario distilled in §3.3.3, cleartext attributes leak information, and single-show credentials reach their limit.

The obvious way to overcome to the single-show limitation when using Brands' credentials is to obtain a number of signed credentials at the time of issuance and use each of these "single redemption coupons" credentials once and only once. One coupon is used each time a show is required to any verifier. The anonymous credentials scheme of Camenisch and Lysyanskaya offers the possibility of multiple show and could be used instead. Due to the construction of anonymous credentials, neither the cleartext attribute problem nor the single-show problem in Brands' system is present.

Adaptability. With the decoupling of matchers and thresholds, with the omission of fuzzy extractors, *Design 1* provides better adaptability. *Design 1* allows each authority to choose matchers and thresholds, whereas *Design 2* requires agreement across the community regarding fuzzy extractor parameters, which naturally includes the error-correcting distance. As mentioned above, this feature of *Design 2* can be problematic in that it becomes impossible for verifiers to decouple their threshold choices from the larger community. To appreciate the value of this adaptability, it suffices to consider the need for an operational verifier to change match thresholds due to environmental variances in the field (such as amber alerts) or accommodate temporary over-illumination or variances in the level of image quality between e-passport issuances.

7.3 Access Control

This section discusses a strategy-based pattern for checkpoint adaptivity approach which compliments the attribute based adaptivity presented in chapter 5. As well, we relax the initial assumptions on discrete-time sampling intervals and fixed zone perimeters. Finally, we discuss how rbCAC and the familiar XACML architecture work

together, and the applicability of rbCAC within a single enterprise as opposed to in an ecosystem.

7.3.1 Strategy-based Adaptivity

Chapter 6 describes a fine-grained adaptive interviewing algorithm. A coarse-grained "strategy-based" pattern of adaptivity is also possible (Gamma et al., 1993). At the checkpoint level, an authority may have different ways to process a transaction, not all of which are equal in terms of cost or quality of output. Figure 29 demonstrates a sample strategy pattern for the PreboardCheckpoint. The cyber-physical interchange between the PreboardCheckpoint and UserAgent is shown. During verification, three alternative strategies are available to the PreboardCheckpoint. For low-risk travellers, a fully automated strategy is available. Under a second strategy, a remote attendant to interview travellers is also available, which can be assumed to bring a slightly higher transaction cost but also a higher information level of assurance regarding post-conditions. A third strategy, in which an in-person interview supplements the device-to-device interchange, is also available for higher-risk situations.

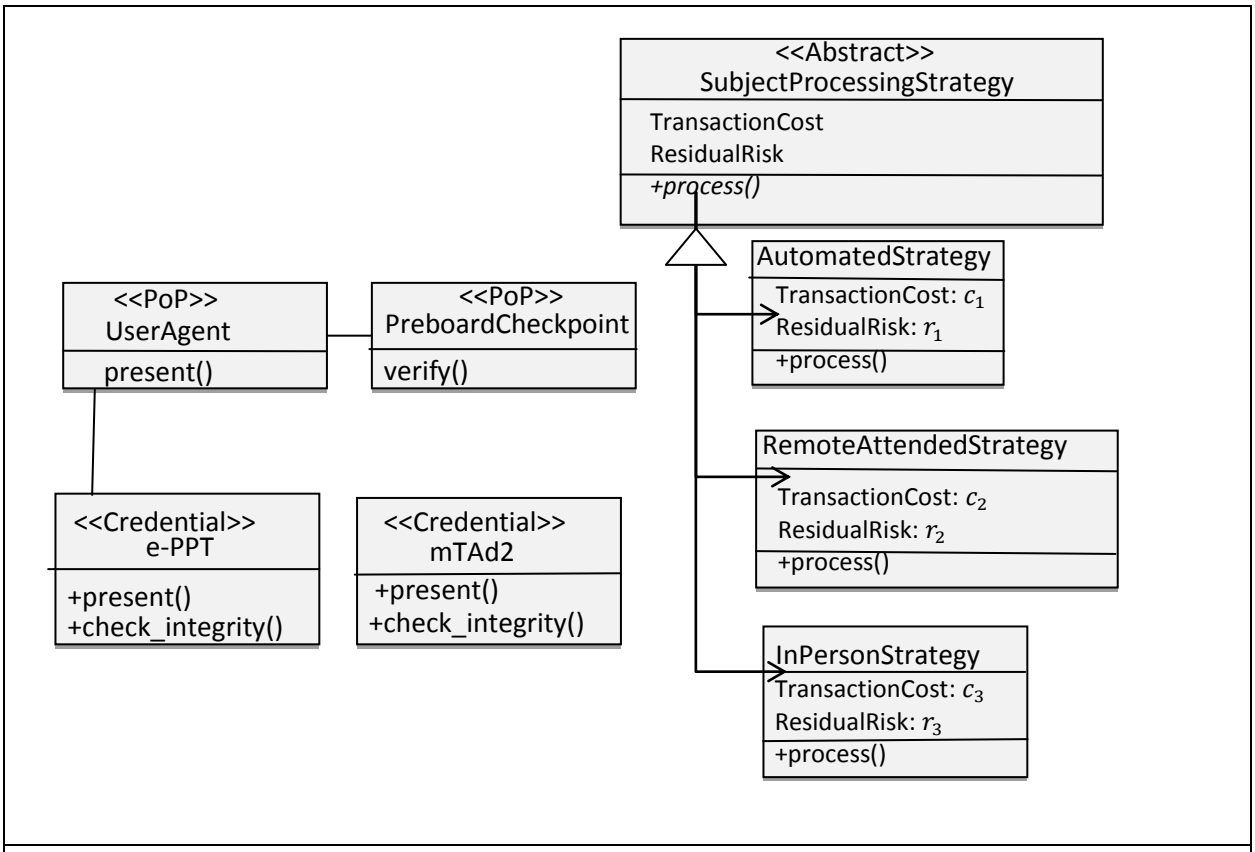

Figure 29) Strategy-based Checkpoint Adaptivity with Sample Contracts

The authority selects the processing strategy according to the specifics of the situation and the subject at hand. The strategy selected has an impact on local and global performance objectives. A low-cost screening strategy may be accorded to low-risk situations. Another strategy may be more costly but have a lower error rate and thus produce a better input for downstream checkpoints.

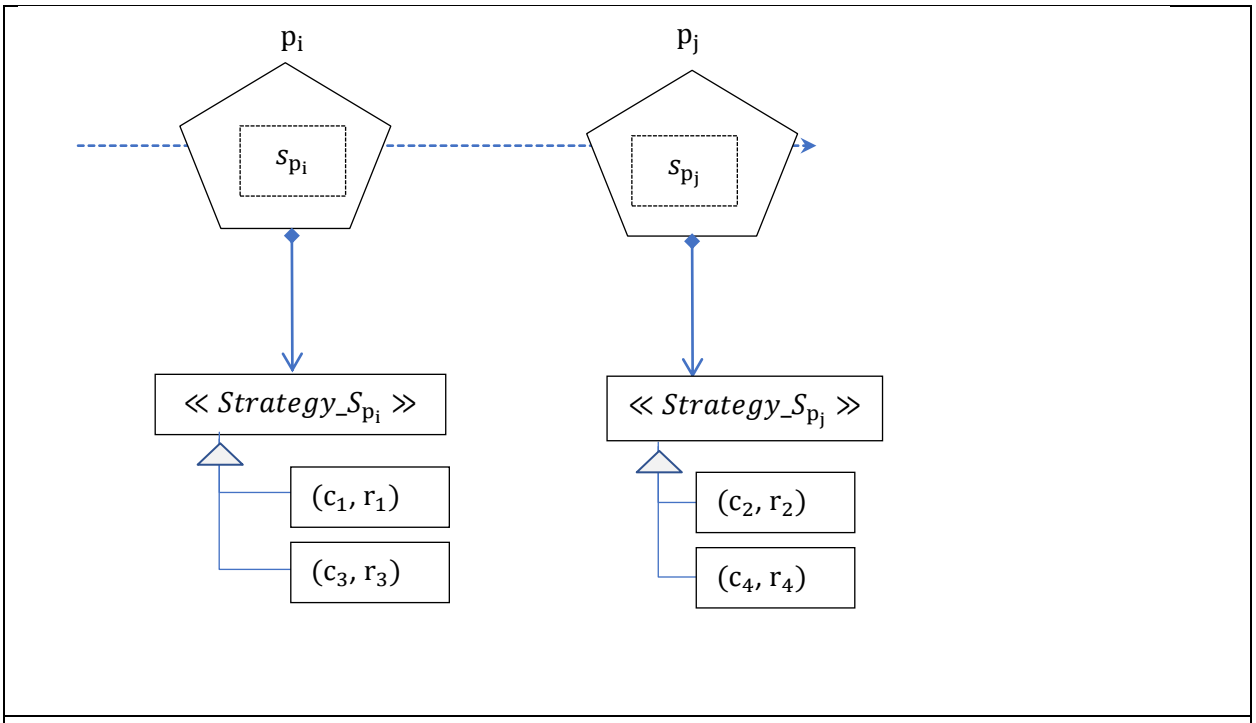

Figure 30) Screening Strategies and Pipeline Effect

Figure 30 shows the components involved in a simplified strategy pipeline between two PoPs in an rbCAC concourse. The checkpoints form a pipeline, and each checkpoint has two alternative strategies. The subtypes of a checkpoint's strategies may differ in many respects (expected interview time, personnel and equipment cost, perceived obtrusiveness, assurance level, residual risk, expected safety, etc.). Predicates may be important enough to be published as a design contract on a class of strategy. Expected values of FPR, TPR, and cost can be treated in this manner.

Let us assume costs and residual risks $\{(c_i, r_i)\}$ where $i \in [1,4]$. Let $c_1 \leq c_2 \leq c_3 \leq c_4$ and $r_1 \geq r_2 \geq r_3 \geq r_4$. In this situation, the cost is inversely proportional to the residual risk. The more a checkpoint interview costs, the less residual risk is left with respect to the screened subject. If provider 1 chooses a low-cost strategy, the residual

risk which must be mitigated by provider 2 is larger. Provider 2 may need to impose stricter screening to protect downstream providers. Any risky subject that evades interdiction due to the screening processes of P1 may inflict damage on the protected resource (which, in the case of a pre-boarding checkpoint, is the airplane).

Choosing a higher-cost strategy upstream may increase the local and global costs; however, it may benefit downstream authorities by removing risk from their doorstep. It may also benefit the overall ecosystem owner by removing the impact of an adversarial subject early in the transaction lifecycle.

7.3.2 Continuous Time and Dynamic Contours

Two implicit assumptions were made for rbCAC on the ATBS setting in Chapter 6: first, that the sampling of subject data occurs at discrete time intervals (i.e., when a subject explicitly reports seeking access to a ROI); and second, that ROI has fixed contours. In many situations, tighter sampling times towards a real-time "instantaneous" risk assessment may be more effective. Such situations include, for example, the monitoring of SCADA equipment or the detection of risk of insider threats. Similarly, certain situations are better represented using dynamic contours of $Regions\ of\ Interest$. Here, for example, we can consider the containment of forest fires, oil spills, or virus epidemics. For ATBS, discrete time and contour assumptions may have been a realistic starting point. We examine the relaxation of these assumptions.

The time interval at which the subject or environment is queried may be quite large or, indeed, quite small. The set of observations collected over time can thus be seen as

a time series over either a discrete- or continuous-time random variable, depending on the granularity of the effective sampling rate.

For coarse-grained sample rates, calculations such as average perceived risk are based on sums and differences. If the sampling rate of subject attributes is decreased such that it becomes infinitesimally short, the same calculations use derivatives. The approximation becomes more precise, and its value at any given time becomes more meaningful as the sampling time shrinks.

In a fixed contour setting, the rate of change of the contour of an ROI and the probability that a contour may have changed over a sampling interval is zero. When the contour of an ROI is allowed to change, the probability that it will have changed between observations is non-zero. Depending on the rate of change of the contour and the size of the sampling interval, the change in the contour between observations may be quite large. Contour changes over time affect various risks. When contours are fluid, a continuous sampling rate yields more confident estimates of current size and the various risk calculation and routing characteristics that accompany the model. Continuous sampling and dynamic ROIs can be applied to the nested and aggregate ROI model presented in this thesis.

7.3.3 rbCAC and XACML

Figure 31 presents rbCAC in terms of the XACML architecture. Each checkpoint is given an infrastructure akin to the XACML architecture. The stacks are connected into a graph using an appropriate concourse mechanism, after which the resulting network is

connected to a stream of subjects to be processed. Each PoP is supported by a stack in which the corresponding architectural components perform a similar role as in traditional XACML.

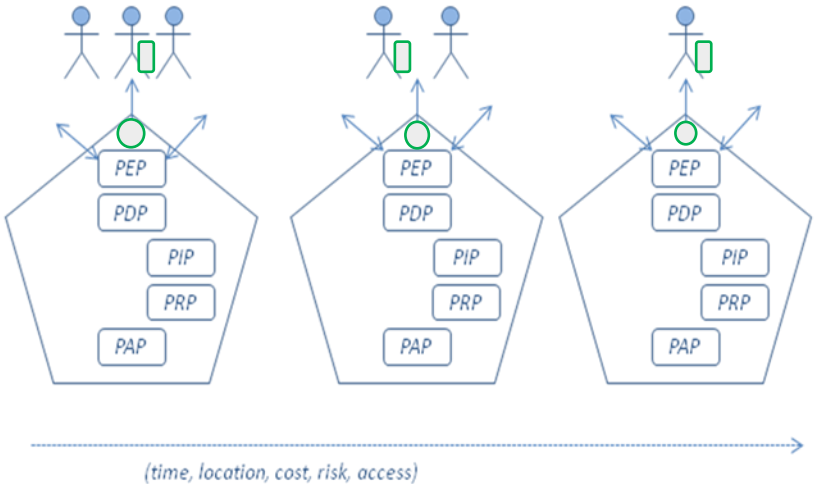

Figure 31) XACML architecture for rbCAC

7.3.3.1 Point of Presence

The *PoP* is not part of XACML. In XACML, the existence of interface/façade systems is largely implied. In EoS-UML and in *rbCAC* , the *PoP* represents the stakeholders' view onto and control of the environment and the transactions which are conducted to generate value. As discussed in §3.3.1, all transacting stakeholders have computer systems that act as their PoP, represent their interests, and give a view of current and past transactions. Context information may include information from environment

sensors, upstream systems, or a community transcript. The acquired subject data consists of the attributes collected about the subject of the transaction during the request-response interview process. As in §6.3.3.2, this interview may be dynamic and risk-aware.

7.3.3.2 Policy Enforcement Point

The policy enforcement point (PEP) is an enterprise component which controls access to a protected area. It works similarly in rbCAC but may take on an explicit cyber-physical role. In CPE and rbCAC , the PEP interfaces with the authority PoP and also includes actuators as part of routing control. The authority PEP may interface with different PoPs depending on the exception processing requirements. The PEP provides the environment actuation interface for the node. Each node communicates with the environment in three manners: acquisition of context information, acquisition of subject data, and decision-based actuation. The first two are communications between $PoPs$, while the third is strictly a PEP function. "Actuation" consists of the node communicating to the subject and the environmental aspects of the decision it has made. This actuation can include writing to a shared ledger, communicating instructions on a PoP user interface, and actuating environment peripheral hardware (such as lighting and barrier gates).

7.3.3.3 Policy Decision Point

The Policy Decision Point (PDP) is a core algorithm which is distinct to each authority but may be shared by the PoP/PEP of a single authority. The PDP contains a service access interface for core algorithms. Thresholds are coordinated at the PDP, which may

invoke exception processing. The PDP is responsible for executing the appropriate policy for the transaction at hand, including interfacing for policy retrieval, activating interview acquisition channels, providing decisions and instructions to PEP which sequences acquisitions, responses, and activations. The PDP evaluates interview responses and environment data to determine applicable risk, mitigations, and the sequencing of appropriate response actions. The $policy\_stabilize\_risk()$ algorithm of Figure 21, above, is a PDP function.

7.3.3.4 Policy Information Point

As in XACML, the Policy Information Point (PIP) includes internal systems and oracle services in the environment. This may include access to the ecosystem directory, transcript, and external “Oracle” services.

In Figure 31, above, each node is the full XACML stack. In general, however, it is more likely that an authority has more than one PoP and shares backend services for them all. If more than one node on the concourse is operated by the same authority, it is possible to share backend services, providing different, specific PoPs for the various contexts in which the authority operates.

Applying the stack to an authority in ATBS, the airline might, for example, have a multi-lane PEP for boarding control, in which each lane has a physical camera and NFC reader for input, and a turnstile and overhead display for output. The airline presents a web-delivery proximity-sensitive PoP to the secure browser on a user’s smartphone. The browser app collects credential, biometric, and travel data as needed and submits it to the airline’s back-end PDP for risk-based decision-making given the auxiliary data

collected by the PIP back-end service. The airline PIP might a) interface with the transcript to get risk assessment results produced and shared by the immigration authority; b) access directory services for public keying information; c) use a face comparison service; or d) make an online call to the board/no-board service of the border authority. PEP activation might include dispensing a “boarding-pass” crypto-credential, presenting appropriate display instructions, and activating a green-light, turnstile-unlock, and photo-capture event. If human interaction were required, a notification could be sent to an employee PoP, who would be charged with escalated processing.

7.3.4 rbCAC within an Enterprise

Chapter 6 elaborates and demonstrates rbCAC in a multi-authority scenario; however, it can be used within an enterprise, with components distributed between internal departments. Consider, for example, a procurement scenario in a large government agency in which requests for proposals and proposal submissions are the subjects of interest. As another example, consider a hiring scenario in a large, geographically distributed digital ecommerce giant, in which candidate screening is the theme. In these scenarios, the various collaborative authorities and systems of the ecosystem become the specialized departments and internal processes they use. Trust increases in these scenarios. This increased inter-authority trust ripples throughout the model. The hierarchical objectives, and the possible contradictions between them, become increasingly synergistic, moving from the Stackelberg model in game theory to the more common problem of optimization of a single (albeit complex) objective

function. The hierarchical objectives can be expected to be more aligned between local and global objectives, with less contradiction/trade-off required. The requirements on the community transcript become simpler. Since trust between authorities is higher, a traditional database rather than a distributed ledger can be used. The benefit of strong contracts remains clear. More tolerance might be allowable, and better optimization may be afforded. As this trust increases and a single focus of objective emerges, we move away from the ecosystem and toward the traditional system (albeit possibly large and geographically distributed).

7.4 Chapter Conclusion

This chapter discusses selected topics in in EoS-UML, Credential Design and rbCAC and the ecosystem as a distributed collaborative processing platform. Ecosystem topics include DbSC, other types of ecosystem, CPE-TM and attack scripting.
DbSC features ledger resident pre-conditions, post-conditions and invariants for transactions and points of presence in the environment, between which stakeholders conduct transactions. Other types of ecosystem may be useful as ecosystem analysis develops. Axes for taxonomy include the type of exchange (product or service) or characteristics of Regions of Interest (static vs. dynamic contours), or characteristics of the subject environment time sampling interval (discrete-time vs. continuous time sampling) of the nodes or of the entire ecosystem.

This Chapter also summarized the differences in terms of the candidate EbC and AbC (Design 1 and Design 2 resp.) were presented in Chapters 4 and 5. Tamper-resistance and unforgeability are delivered by the properties of the signature schemes

used for the credentials. Design 1 allows for standard cryptography to be used and permits easier adoption with IT organizations. The privacy features of Design 2 require well-studied but non-standard cryptography. Privacy and Operational properties become the large differentiators with Design 1 offering no selective show or credential composition and Design 2 providing the range of required properties. Operational properties become more challenging for Design 2 for a number of reasons. The use of non-standardized digital signatures can make adoption more difficult for government agencies. Design 1 becomes more attractive to such organizations. A key security property of focus is non-lendability. Fuzzy extractor algorithms were used for non-lendability in each case. Field trials in Design 1 noted performance differences between traditional matchers and fuzzy extractors. Analysis also shows a challenge on decoupling the verifier and issuer error correction thresholds. These two considerations significantly impact operational qualities required of Design 2 in a multi-stakeholder SoS setting. Biometric performance can be obtained using fit for purpose biometric matchers but requires assumptions of Honest-but-Curious service providers. These assumptions may or may not be realistic considering governments as opposed to corporations, and also differences in international government privacy policies.

Topics on rbCAC include a strategy-based pattern for checkpoint adaptivity approach which complements the attribute-based adaptivity, a discussion on discrete-time sampling intervals and fixed zone perimeters, a discussion of rbCAC within a single-enterprise and a discussion on the synergies between rbCAC and XACML.

Chapter 8 Conclusion and Future Work

8.1 Conclusion

This thesis proposed CPE as an enhancement to CPS and further described an EoS-UML. The EoS-UML allows for definitions of a novel CPS threat model. At the heart of the ecosystem is the stakeholder mandate to drive benefit. The assurances required to mitigate risk to benefit are delivered by cryptographic credentials, which fall into one of two broad types: envelop-based and attribute-based. These two varieties provide different benefits to stakeholders. Often there are trade-offs between properties, for example, with a “selective show” privacy feature often coming at an operational cost in terms of “complexity of integration” due to non-standard cryptography. Non-lendability is a common property required in the digital world. Both sample credential designs used a fuzzy extractor pattern to provide biometrically supported non-lendability. The important distinction between the performance and measurement of fit-for-purpose biometric-matching algorithms and fuzzy extractors was discussed in the present work. Assuming a performant credential and design contracts on stakeholders who accord each other various levels of (dis)trust, collaborative-distributed computing is possible. A novel access-control mechanism, rbCAC , was also presented in this research. The ecosystem as a platform for computing has certain analogies to artificial intelligence (hierarchical multi-objective optimization and neural computing) that may be explored in future work.

8.1.1 CPE and EoS UML

The thesis proposed an extension of CPS that features multiple actors in an environment in which human subjects are an implicit part of the systems loop. CPE is a cyber–physical system in which subjects interact with multiple authorities and every stakeholder must weigh the potential costs and benefits of actions with respect to their own objectives.

To our knowledge, no such formalism exists in the CPS literature. We applied this formalism on a CPS of growing importance: air travel and border security. To accompany the CPE concept, we proposed a threat model that similarly included a cyber-physical attacker with human and machine elements that had the ability to conduct distributed attacks through the CPS credential.

This CPE formalism is an extension of the CPSS and CPS paradigms. The multi-actor perspective, goal-driven activities, and cost/benefit decisions that take place within CPE could be useful to analysts and designers of CPS and CPSS in designing terrain checkpoints and ecosystems. The same formalism could potentially help system engineering on large scale, multi-authority systems, such as government food stamp programs, citizen health and safety applications, and applications of international air travel and border security. The formalism could also be used across a wide range of cyber physical systems, particularly those that require the sensitive treatment of uncertain data and make decisions pertaining to citizen wellbeing.

The threat model we proposed featured both cyber and physical components and was distributed so that the threat and the CPS terrain were tightly integrated. We feel that this represents a valuable contribution to the body of knowledge on CPS.

8.1.2 CPE Threat Model

The proposed threat model builds on the standard threat model used in cryptography by mirroring characteristics of CPE and incorporating human, geospatial, and temporal elements. The proposed CPE attacker $\widetilde{\mathrm{A}}$ is composed of three components: the command-and-control attacker $\widetilde{\mathrm{Ac}}$, the terrain attack force $\widetilde{\mathrm{At}}$, and the attack platform $\widetilde{\mathrm{Ap}}$, $\widetilde{\mathrm{A}} = (\widetilde{\mathrm{Ac}}, \widetilde{\mathrm{At}}, , \widetilde{\mathrm{Ap}})$. The $\widetilde{\mathrm{Ac}}$ and $\widetilde{\mathrm{At}}$ are humans equipped with compute devices. $\widetilde{\mathrm{Ac}}$ acts in the Command-and-Control role. $\widetilde{\mathrm{At}}$ is the attack ground force. Mobile within the terrain, $\widetilde{\mathrm{At}}$ performs intelligence, surveillance, reconnaissance, and engagement tasks. $\widetilde{\mathrm{Ap}}$ is the algorithmic platform required for the attack.

$\widetilde{\mathrm{Ac}}$, $\widetilde{\mathrm{Ap}}$, and $\widetilde{\mathrm{At}}$ are instantiated at the start of an attack with objectives, algorithms, and personnel. The size of $\widetilde{\mathrm{At}}$ changes over time. $\widetilde{\mathrm{At}}$ grows as honest actors in the field (or their computer equipment) are corrupted. $\widetilde{\mathrm{Ap}}$ grows as $\widetilde{\mathrm{A}}$ conquers core systems and databases. Expansion allows $\widetilde{\mathrm{A}}$ to gain a deeper penetration into ecosystem contexts and access equipment, protected data, and additional algorithms. Context penetration can provide a richer overall vantage point and ecosystem attack surface.

8.1.3 Comparative Credentials Design

We proposed the credential as a vehicle for value-interchange in a CPE. Credentials allow transactors to conduct electronic transactions in a manner that balances security, privacy, and operational concerns. We proposed a basis set of properties, knowing that others will emerge, depending on stakeholder requirements.

We demonstrated the trade-offs that alternative credential designs may exhibit with respect to the goals of different ecosystem stakeholders. We proposed two designs, both of which achieved non-lendability through privacy-respecting biometric verification. Both had different characteristics considering subject or privacy sensitivity and strong security and interoperability. One design used a fuzzy extractor secret as a biometrically derived key to secure a symmetrically encrypted, asymmetrically signed credential, illustrating how target mobile credential algorithms can be implemented in today's environment. Meanwhile, the other design used derived key generation as an embedded secret in a Brands' digital credential scheme and met additional privacy-by-design requirements. Both algorithms are novel, as is their application and demonstration in the ATBS domain. These algorithms support a key premise of this thesis, which is that when engineering ecosystems, the design of the digital credential vehicle that is used in transactional risk analysis is fundamental, and in credential design, a given design may not fit the desires of all stakeholders.

8.1.4 Risk-Balanced Cellular Access Control

We presented a distributed multi-authority, risk-aware, decision-making model and applied it to the problem of subject screening and access control faced with deception. To our knowledge, this has not been conducted before. Previous risk-aware access-control approaches have considered risk-balanced access control in a single authority/ perimeter setting. In our setting, the zone perimeter was composable recursively and as an aggregate. Screening and access control thus occurred not only once but repeatedly throughout a subject’s interactions in the ecosystem. The progressive nature of our

approach was also new. As multiple stakeholders are involved in processing a subject, the risk-assessments of one authority can benefit the screening processes of subsequent authorities downstream. This brings us toward a continuous risk assessment and access-control pattern in contrast to the traditional single-perimeter approach.

rbCAC can be applied within a single enterprise in a manner similar to a multi-authority implementation. For this, an enterprise coordination function can be substituted for the overarching authority, and there can be separate departments for local authorities. The example of human-resource screening in a multinational organization can be used to illustrate rbCAC within enterprises, while to examine rbCAC in an ecosystem and within the enterprise, the example of a home purchase and subsequent mortgage approval may be used.

8.2 Future Work

There is scope for significant future research and business opportunities in this area. These are discussed in terms of the CPE conceptual and threat models, credentials, and distributed processing and intelligent systems.

8.2.1 CPE Threat Model

Future work could present a case study demonstrating the usage of the requirements analysis and design artefacts presented in this thesis.

8.2.2 EoS-UML

The conceptual model of the CPE could be augmented with pre- and post-conditions and invariants. Such work should be combined with probabilistic correctness in terms of statistical Type 1 and Type 2 errors. The properties relevant to privacy-respecting work in crypto-currencies, mix-nets, and the onion-router should be outlined, including those with scientific, operational, and social impacts. The property review must also remain pragmatic and operationally meaningful, such as through including the CPS domains of healthcare, naval and airspace protection, and food stamps.

Future research should also seek to further refine EoS-UML and CPE-TM to include continuous time sampling, dynamic contours, and non-collaborative (competitive/adversarial) trust models.

8.2.3 CPE Threat Model

The threat model and attacker of CPE presents a starting point for recognizing the highly pertinent concepts of objectives, location, human capability, and timing, which have not received adequate focus in traditional security analysis and CPS applications. A focused application of CPE-TM on a relevant CPS or SCADA attack could provide interesting results, and Stuxnet could represent a good starting point. A survey of CPS applications and threat models where ecosystem-thinking and CPE-TM may apply should be conducted.

8.2.4 Credential Design

This thesis offers a starting point for discussing comparative applied credential design. Credential feature enhancements in performance, non-lendability, auditability, and accountability will continue to be of interest to the scientific and operational communities as we expect ecosystem properties and requirements to evolve as stakeholder expectations increase.

8.2.4.1 Operational Fuzzy Extractors

Fuzzy extractors may benefit from increased configurability in operations, which could be achieved by focusing on ROC curve analysis and decoupling issuer and verifier thresholds.

8.2.4.2 Credential enhancements

Depending on user requirements, a number of enhancements can be made to AbC and their possible combinations using decentralized ledger technologies, including homomorphic encryption for closed-set annotations, accumulators and zero-knowledge proofs for credential issuance and revocation, and commitments for proofs of related transactions. Naturally, enhancements to the usability and performance of privacy-respecting, non-transferability techniques will continue to provide value.

8.2.5 Distributed Processing Model and Architecture

The proposed architecture and processing model has a mixture of interesting characteristics from a theoretical and physical perspective.

8.2.5.1 Risk-balanced Cellular Access Control

Future work should include research into optimize configuration for screening queries and risk thresholds in rbCAC. This should include simulated annealing and genetic algorithms approaches for find optimal solutions subject to hierarchical objective structure of global and local stakeholders. This can lead to an adaptive collaborative screening platform.

8.2.5.2 Turing Machine and Automata Theory

The state machine representation in Chapter 7 recalls automata theory. Considering the non-deterministic behavior inherent in the SoS component model places us in the domain of non-deterministic finite state automata. Here, Markov models could be used to model the system, while the model could be categorized within automata, and Turing machine theory could be applied in future work.

8.2.5.3 Architecture

Future work should include defining and categorizing the collaborative architecture outlined in §6.3.1, particularly with the decentralized PoP and the Design by Smart Contract approach method for solidifying trust. Our architecture uses has aspects of the component and distributed architecture, shares features with the Von Neumann Architecture, the Harvard Architecture, and Blockchain Architecture models.

From a deployment perspective, the model has commonalities with IoT/Fog architectures, as well as Cloud and traditional “Hub-and-Spoke” models. CPE should be studied in conjunction with Fog computing from application and security perspectives.

8.2.5.4 Full-stack security and privacy

Today's CPE environments are made up of technologies that span decades of technological legacy which include host transaction systems, client-server approaches, web-delivery presentation and server methods, cloud computing, and consumer smartphone operating systems. These components have varying level of vulnerabilities and include programmer bugs as well as vendor-introduced backdoors, and deceptions. This varied assortment of technologies has not been certified to offer the data security and privacy that government-to-citizen systems require. Stakeholder objectives, including those of international hardware providers, digital giant platforms, search engine providers, and infotainment apps, should also be included in threat analysis, and at the basis of operational assumptions.

8.2.6 Optimization and Intelligent Systems

8.2.6.1 Hierarchical Objective Optimization

Future work should apply an rbCAC questionnaire and threshold optimization to maximize ecosystem global objectives while achieving a satisfactory balance across a possibly conflicting set of objectives between local authorities.

8.2.6.2 Neural Network Analogy

The distributed architecture and rbCAC presented in Chapter 6 have certain parallels to the conceptual model used in neural computing. First, they feature a collection of distributed connected processing elements with thresholded activation. The parallel to neural computing is in the multiple connected neurons and thresholded sigmoid activation function. Further, each node can implement its specific input processing

algorithm. In addition, a high-level processing task, as defined in rbCAC , involves classifying input vectors into risk-profile/activation functions. This mirrors a standard neural net task of categorizing input (e.g., faces or hand-written text) into known categories of output. Finally, calibrating an rbCAC network instance consists of tuning the screening questionnaires and numeric thresholds of its nodes. This parallels the learning function in a neural network, where iterative refinement or backpropagation of erroneous classification is used to train the network. These factors suggest that an ecosystem computation network can be used as a distributed neural net or classifier to process incoming observations into appropriate outputs

References

ACR122U USB NFC Reader, http://www.acs.com.hk/en/products/3/acr122u-usb-nfc-reader/, last accessed 2018/09/12.

Adams, C. (2011). Achieving non‐transferability in credential systems using hidden biometrics. Security and Communication Networks, 4(2), 195-206.

Adler, A. (2004). Images can be Regenerated from Quantized Biometric Match Score Data, Proceedings of the Canadian Conference on Electrical and Computer Engineering, Niagara Falls, Canada, 469-472

Ahonen, T., Hadid, A., Pietikainen, M.: 'Face description with local binary patterns: application to face recognition', IEEE Trans. Pattern Anal. Mach. Intell., 2006, 28, (12), pp. 2037–2041

Altman, D. G., & Bland, J. M. (1994). Diagnostic tests. 1: Sensitivity and specificity. BMJ: British Medical Journal, 308(6943), 1552.

Altman, D. G., & Bland, J. M. (1994). Statistics Notes: Diagnostic tests 2: predictive values. Bmj, 309(6947), 102. HTTPS://www.ncbi.nlm.nih.gov/pmc/articles/PMC2540558/pdf/bmj00448-0038a.pdf

Anda, Amal Ahmed, and Daniel Amyot. "Arithmetic semantics of feature and goal models for adaptive cyber-physical systems." *2019 IEEE 27th International Requirements Engineering Conference (RE)*. IEEE, 2019.

Atlam, Hany, et al. "An overview of risk estimation techniques in risk-based access control for the internet of things." (2017): 254-260.

Atlam, Hany F., et al. "Risk-Based Access Control Model: A Systematic Literature Review." Future Internet 12.6 (2020): 103.

Axelrod, C. Warren. "Managing the risks of cyber-physical systems." *2013 IEEE Long Island Systems, Applications and Technology Conference (LISAT)*. IEEE, 2013.

Bagnato, Alessandra, et al. "The INTO-CPS cyber-physical system profile." *Ada User* 38.4 (2017): 227.

Baldimtsi, F., & Lysyanskaya, A. (2013, December). On the security of one-witness blind signature schemes. In International Conference on the Theory and Application of Cryptology and Information Security (pp. 82-99). Springer Berlin Heidelberg.

Baldimtsi, F., & Lysyanskaya, A. (2013, November). Anonymous credentials light. In Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security (pp. 1087-1098). ACM

Banerjea-Brodeur, Nicolas Paul. Advance Passenger Information/Passenger Name Record: Privacy Rights and Security Awareness. Diss. McGill University, Montreal, 2003.

Bartlett, M., Movellan, J., Sejnowski, T.: 'Face recognition by independent component analysis', IEEE Trans. Neural Netw., 2002, 13, (6), pp. 1450–1464

Basar, Tamer, and Jose B. Cruz Jr. Concepts and Methods in Multi-Person Coordination and Control. ILLINOIS UNIV AT URBANA DECISION AND CONTROL LAB, 1981.

Başar, Tamer, and Geert Jan Olsder. Dynamic noncooperative game theory. Society for Industrial and Applied Mathematics, 1998.

Battram, Peter, Bernhard Kaiser, and Raphael Weber. "A Modular Safety Assurance Method considering Multi-Aspect Contracts during Cyber Physical System Design." *REFSQ Workshops*. 2015.

Bichsel, P., Camenisch, J., Groß, T., & Shoup, V. (2009, November). Anonymous credentials on a standard java card. In Proceedings of the 16th ACM conference on Computer and communications security (pp. 600-610). ACM.

Bissessar, D. (2013). Cryptographic credentials with privacy-preserving biometric bindings (Dissertation, Université d'Ottawa/University of Ottawa).

Bissessar, D., Adams, C., & Liu, D. (2014, July). Using biometric key commitments to prevent unauthorized lending of cryptographic credentials. In Privacy, Security and Trust (PST), 2014 Twelfth Annual International Conference on (pp. 75-83). IEEE.

Bissessar, David, et al. "Mobile Travel Credentials." *International Symposium on Foundations and Practice of Security*. Springer, Cham, 2018.

Bitansky, Nir, et al. "Why “Fiat-Shamir for proofs” lacks a proof." Theory of cryptography conference. Springer, Berlin, Heidelberg, 2013.

Blanton, M., & Aliasgari, M. (2011, July). On the (non-) reusability of fuzzy sketches and extractors and security in the computational setting. In Security and Cryptography (SECRYPT), 2011 Proceedings of the International Conference on (pp. 68-77). IEEE.

Blanz, V., Vetter, T.: ‘Face recognition based on fitting a 3D Morphable model’, IEEE Trans. Pattern Anal. Mach. Intell., 2003, 25, (9), pp. 1063–1074

Bogari, Eyad Abdullah, et al. "An investigative analysis of the security weaknesses in the evolution of RFID enabled passport." International Journal of Internet Technology and Secured Transactions 4.4 (2012): 290-311.

Bonci, Andrea, et al. "RMAS: relational multiagent system for CPS prototyping and programming." *2018 14th IEEE/ASME International Conference on Mechatronic and Embedded Systems and Applications (MESA)*. IEEE, 2018.

Booch, G., Rumbaugh, J., and Jacobson, I. The Unified Modeling Language User Guide. Addison-Wesley, Reading, MA, 1999.

Boyen, Xavier. "Reusable cryptographic fuzzy extractors." In Proceedings of the 11th ACM conference on Computer and communications security, pp. 82-91. ACM, 2004.

Brands, S. A. (2000). Rethinking public key infrastructures and digital certificates: building in privacy. MIT Press.

Brien, Renaud. Security, Privacy and Performance Improvements for Fuzzy Extractors. Diss. Université d'Ottawa/University of Ottawa, 2020.

Bringer, J., & Chabanne, H. (2008, June). An authentication protocol with encrypted biometric data. In International Conference on Cryptology in Africa (pp. 109-124). Springer Berlin Heidelberg.

Burmester, Mike, Emmanouil Magkos, and Vassilis Chrissikopoulos. "Modeling security in cyber–physical systems." *International journal of critical infrastructure protection* 5.3-4 (2012): 118-126.

Camenisch, J., & Lysyanskaya, A. (2001, May). An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In International Conference on the Theory and Applications of Cryptographic Techniques (pp. 93-118). Springer Berlin Heidelberg.

Camenisch, J., & Lysyanskaya, A. (2002, September). A signature scheme with efficient protocols. In International Conference on Security in Communication Networks (pp. 268-289). Springer Berlin Heidelberg.

Camenisch, J., & Lysyanskaya, A. (2004, August). Signature schemes and anonymous credentials from bilinear maps. In Annual International Cryptology Conference (pp. 56-72). Springer Berlin Heidelberg.

Canard, S., Lescuyer, R., & Traoré, J. (2011, December). Multi-show anonymous credentials with encrypted attributes in the standard model. In International Conference on Cryptology and Network Security (pp. 194-213). Springer Berlin Heidelberg.

Cao, Yan, et al. "A topology and risk-aware access control framework for cyber-physical space." *Frontiers of Computer Science* 14.4 (2020): 1-16.

Cavoukian A, Stoianov A: Biometric encryption: the new breed of untraceable biometrics. In Biometrics: Fundamentals, Theory, and Systems. Wiley, London; 2009.

CBP ESTA Webpage, http://www.cbp.gov/travel/international-visitors/esta(accessed March 28, 2014)

Chaum, David. "Blind signatures for untraceable payments." In Advances in Cryptology: Proceedings of Crypto, vol. 82, pp. 199-203. 1982.

Chaum, David. "Security without identification: Transaction systems to make big brother obsolete." Communications of the ACM 28, no. 10 (1985): 1030-1044.

Chen, Peter Pin-Shan. "The entity-relationship model—toward a unified view of data." *ACM transactions on database systems (TODS)* 1.1 (1976): 9-36.

Chegini, Hossein, et al. "Process Automation in an IoT–Fog–Cloud Ecosystem: A Survey and Taxonomy." *IoT* 2.1 (2021): 92-118.

CIC eTA Webpage, http://www.cic.gc.ca/english/department/acts-regulations/forward-regulatory-plan/eta.asp (accessed February 27, 2015)

Cole, Eric. Advanced persistent threat: understanding the danger and how to protect your organization. Newnes, 2012.

Cloutier, Rob, et al. "Modeling a system of systems using UML." *Proceedingsof the Conference on Systems Engineering Research, Hoboken, NJ*. 2003.

COURNOT, A. A. 1838. Recherches sur les Principes Mathematiques de la Théorie des Richesses. In Libraire de Sciences Politiques et Sociale. M. Riviere & Cie, Paris.

Cruz Jr, J. B. "Survey of Nash and Stackelberg Equilibrim Strategies in Dynamic Games." *Annals of Economic and Social Measurement, Volume 4, number 2*. NBER, 1975. 339-344.

Daemen, J., Rijmen, V.: The design of Rijndael. Springer-Verlag, Heidelberg (2002).

Davida, G. I., Frankel, Y., & Matt, B. J. (1998, May). On enabling secure applications through off-line biometric identification. In Security and Privacy, 1998. Proceedings. 1998 IEEE Symposium on (pp. 148-157). IEEE.

de Champeaux, Dennis. "Object-oriented analysis and top-down software development." *European Conference on Object-Oriented Programming*. Springer, Berlin, Heidelberg, 1991

Degenhardt, Teresa, and Mike Bourne. "When risks meet: The dance of experience, professional expertise and science in border security technology development." Criminology & Criminal Justice 20.2 (2020): 207-225.

Derler, Patricia, et al. "Cyber-physical system design contracts." *Proceedings of the ACM/IEEE 4th International Conference on Cyber-Physical Systems*. 2013.

Dodis, Y., Reyzin, L., & Smith, A. (2004, May). Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In International Conference on the Theory and Applications of Cryptographic Techniques (pp. 523-EU540). Springer Berlin Heidelberg.

Chain of Trust,https://www.cbsa-asfc.gc.ca/agency-agence/reports-rapports/bp2020/2018/trust-confiance-eng.html (accessed Nov 2021)

“FASTER-PrivBio Project Summary Report”,https://publications.gc.ca/site/eng/9.874440/publication.html, May 2017 (accessed Nov 2021)

https://cradpdf.drdc-rddc.gc.ca/PDFS/unc200/p801301_A1b.pdf

Drummond, C., & Holte, R. C. (2000, August). Explicitly representing expected cost: An alternative to ROC representation. In Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining (pp. 198-207). ACM.

Faltemier, Timothy C., Kevin W. Bowyer, and Patrick J. Flynn. "A region ensemble for 3-D face recognition." *IEEE Transactions on Information Forensics and Security* 3.1 (2008): 62-73.

Farber, Stephen C., Robert Costanza, and Matthew A. Wilson. "Economic and ecological concepts for valuing ecosystem services." Ecological economics 41.3 (2002): 375-392.

Fawcett, T. (2006). An introduction to ROC analysis. Pattern recognition letters, 27(8), 861-874.

Feng, Y. C., Pong C. Yuen, and Anil K. Jain. "A hybrid approach for face template protection." SPIE Defense and Security Symposium. International Society for Optics and Photonics, 2008.

Ferrara, Matteo, Annalisa Franco, and Davide Maltoni. "The magic passport." IEEE International Joint Conference on Biometrics. IEEE, 2014.

Ferri, Cèsar, José Hernández-Orallo, and Peter A. Flach. "A coherent interpretation of AUC as a measure of aggregated classification performance." Proceedings of the 28th International Conference on Machine Learning (ICML-11). 2011.

Amos Fiat; Moni Naor (1994). Broadcast encryption*. Proc. Advances in Cryptology – CRYPTO '93 (Extended abstract). Lecture Notes in Computer Science. 773. pp. 480–491.* doi:10.1007/3-540-48329-2_40*.* ISBN 978-3-540-57766-9*.*

Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In CRYPTO, pages 186–194, 1986.

Frontex: Best practice technical guidelines for automated border control (ABC) systems, Research and Development Unit, Warsaw, 2012.

Fuentes-Fernández, Lidia, and Antonio Vallecillo-Moreno. "An introduction to UML profiles." *UML and Model Engineering* 2.6-13 (2004): 72.

Gamma, Erich, et al. "Design patterns: Abstraction and reuse of object-oriented design." *European Conference on Object-Oriented Programming*. Springer, Berlin, Heidelberg, 1993.

Gerdes, J. H., Huang, C. T., & Sharaf, M. A. Incorporating biometrics into veiled certificates: preventing unauthorized use of anonymous certificates. Electronic Commerce Research, 1-28.

Goldwasser, S., & Kalai, Y. T. (2003, October). On the (in) security of the Fiat-Shamir paradigm. In Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on (pp. 102-113). IEEE.

Gorodetsky, Vladimir, et al. "Multi-agent technology for air traffic control and incident management in airport airspace." *Proceedings of the International Workshop Agents in Traffic and Transportation, Portugal*. 2008.

Grother, Patrick J., et al. *Report on the evaluation of 2D still-image face recognition algorithms*. US Department of Commerce, National Institute of Standards and Technology, 2011.

Guajardo, J., Mennink, B., & Schoenmakers, B. (2010, December). Anonymous credential schemes with encrypted attributes. In International Conference on Cryptology and Network Security (pp. 314-333). Springer Berlin Heidelberg.

Guo, Guodong, and Na Zhang. "A survey on deep learning based face recognition." *Computer vision and image understanding* 189 (2019): 102805.

Hand, D. J. (2009). Measuring classifier performance: a coherent alternative to the area under the ROC curve. Machine learning, 77(1), 103-123.

Henshaw, Michael J. "Systems of systems, cyber-physical systems, the internet-of-things… whatever next?." *Insight* 19.3 (2016): 51-54.

Hoepman, Jaap-Henk, et al. "Crossing borders: Security and privacy issues of the european e-passport." Advances in information and computer security (2006): 152-167.

Holland, John H. "Genetic algorithms." *Scientific american* 267.1 (1992): 66-73.

Huang, Linan, and Quanyan Zhu. "Analysis and computation of adaptive defense strategies against advanced persistent threats for cyber-physical systems." International Conference on Decision and Game Theory for Security. Springer, Cham, 2018.

Hussain, MD Muzakkir, and MM S. Beg. "Using vehicles as fog infrastructures for transportation cyber-physical systems (T-CPS): Fog computing for vehicular networks." *International Journal of Software Science and Computational Intelligence (IJSSCI)* 11.1 (2019): 47-69.

Hussein, Dina, et al. "Dynamic social structure of things: A contextual approach in CPSS." *IEEE Internet Computing* 19.3 (2015): 12-20.

Hutchins, Eric M., Michael J. Cloppert, and Rohan M. Amin. "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains." Leading Issues in Information Warfare & Security Research 1.1 (2011): 80.

ICAO Document 9303: Machine readable travel documents - part 1-2. Technical report, International Civil Aviation Organization, 2006, 2006

ICAO "Digital Travel Credentials - Virtual Component Data Structure and PKI mechanisms", Release 1.2, October 2020

Privacy by Design Resolution. Jerusalem: 32nd International Conference of Data Protection and Privacy Commissioners (2010). http://www.ipc.on.ca/site_documents/pbd-resolution.pdf

ISO/IEC CD 18013-5:2019(E) Personal identification — ISO-compliant driving licence — Part 5: Mobile driving licence (mDL) application

ISO International Standard ISO/IEC JTC 1/SC 37 N506, Text of FCD 19794–5, Biometric Data Interchange Formats—Part 5: Face Image Data 2004 [Online]. Available: http://isotc.iso.org.

ISO/IEC 19794–5: 2005/Amd 1:2007, Information Technology—Biometric data Interchange format—Part 5: Face Image Data/Amendment 1: Conditions for Taking Photographs for Face Image Data 2007.

ISO/IEC 19794–5: 2005/Cor 1&2:2008, Information Technology—Biometric Data Interchange Format—Part 5: Face Image Data, 2008, Technical Corrigendum 1&2.

ISO/IEC TR 29794-5:2010, Information Technology—Biometric Sample Quality—Part 5: Face Image Data 2010

Information Technology - Security techniques - Biometric Information Protection, ISO/IEC IS 24745, June 2011

Jackson, Michael C., and Paul Keys. "Towards a system of systems methodologies." Journal of the operational research society 35.6 (1984): 473-486.

Jacobson, Ivar. (1992). Object-oriented software engineering : a use case driven approach. [New York]: ACM Press. pp. 130–133. ISBN 0201544350. OCLC 26132801.

Jain, Anil K., and Umut Uludag. "Hiding biometric data." IEEE transactions on pattern analysis and machine intelligence 25.11 (2003): 1494-1498.

Jain, Anil K., Karthik Nandakumar, and Abhishek Nagar. "Biometric template security." EURASIP Journal on Advances in Signal Processing 2008 (2008): 113.

Japkowicz, N., & Shah, M. (2011). Evaluating learning algorithms: a classification perspective. Cambridge University Press.

Japkowicz, N., & Stephen, S. (2002). The class imbalance problem: A systematic study. Intelligent data analysis, 6(5), 429-449.

Jiang, Bo, Xuegong Zhang, and Tianxi Cai. "Estimating the confidence interval for prediction errors of support vector machine classifiers." *The Journal of Machine Learning Research* 9 (2008): 521-540.

JMRTD: An Open Source Java Implementation of Machine Readable Travel Documents, HTTPS://jmrtd.org/, last accessed 2018/09/12.

Johnson, D., Menezes, A., Vanstone, S.: The elliptic curve digital signature algorithm (ECDSA). International journal of information security 1(1), 36-63 (2001).

Juels, Ari, David Molnar, and David Wagner. "Security and Privacy Issues in E-passports." Security and Privacy for Emerging Areas in Communications Networks, 2005. SecureComm 2005. First International Conference on. IEEE, 2005.

Juels, A., & Sudan, M. (2002). A Fuzzy Vault Scheme Proc. In Intl Symp. Inf. Theory, A Lapidoth, E. Teletar, Eds (p. 408).

Juels, A., & Wattenberg, M. (1999, November). A fuzzy commitment scheme. In Proceedings of the 6th ACM conference on Computer and communications security (pp. 28-36). ACM

Kandala, Savith, Ravi Sandhu, and Venkata Bhamidipati. "An attribute based framework for risk-adaptive access control models." 2011 Sixth International Conference on Availability, Reliability and Security. IEEE, 2011.

Kevenaar, T. A., Schrijen, G. J., van der Veen, M., Akkermans, A. H., & Zuo, F. (2005, October). Face recognition with renewable and privacy preserving binary templates. In Fourth IEEE Workshop on Automatic Identification Advanced Technologies (AutoID'05) (pp. 21-26). IEEE.

Khan, Bilal, Muhammad Khurram Khan, and Khaled S. Alghathbar. "Biometrics and identity management for homeland security applications in Saudi Arabia." African Journal of Business Management 4.15 (2010): 3296-3306.

Khan, S., H. Maoh, and W. ANDERSON. "Simulating the Impacts of RFID Enabled Lanes at the Canada-US Border: An Application to the Ambassador Bridge." TAC 2016: Efficient Transportation-Managing the Demand-2016 Conference and Exhibition of the Transportation Association of Canada. 2016.

Klötzer, C., & Pflaum, A. (2015, October). Cyber-physical systems as the technical foundation for problem solutions in manufacturing, logistics and supply chain management. In Internet of Things (IOT), 2015 5th International Conference on the (pp. 12-19). IEEE.

Koning, Merel, et al. "The ABC of ABC: An analysis of attribute-based credentials in the light of data protection, privacy and identity." (2014).

Lafkih, M., Lacharme, P., Rosenberger, C., Mikram, M., Ghouzali, S., El Haziti, M., & Aboutajdine, D. (2015, August). Vulnerabilities of fuzzy vault schemes using biometric data with traces. In Wireless Communications and Mobile Computing Conference (IWCMC), 2015 International (pp. 822-827). IEEE.

Larman, Craig, and Victor R. Basili. "Iterative and incremental developments. a brief history." *Computer* 36, no. 6 (2003): 47-56.

LaPadula, Leonard J., and D. Elliot Bell. Secure computer systems: A mathematical model. Vol. 2. Technical Report 2547, 1996.

Lee, Kyou Ho, Jeong Hee Hong, and Tag Gon Kim. "System of systems approach to formal modeling of CPS for simulation-based analysis." *Etri Journal* 37.1 (2015): 175-185.

Lee, Edward Ashford, and Sanjit A. Seshia. *Introduction to embedded systems: A cyber-physical systems approach*. Mit Press, 2017.

Lehtonen, Pinja, and Pami Aalto. "Smart and secure borders through automated border control systems in the EU? The views of political stakeholders in the Member States." European Security 26.2 (2017): 207-225.

Lenzini, Gabriele, Sjouke Mauw, and Samir Ouchani. "Security analysis of socio-technical physical systems." *Computers & electrical engineering* 47 (2015): 258-274.

Leveson, Nancy. "Perspective: The Drawbacks in Using the Term 'System of Systems'." *Biomedical instrumentation & technology* 47.2 (2013): 115-118.

Li, Q., Sutcu, Y., & Memon, N. (2006, December). Secure sketch for biometric templates. In International Conference on the Theory and Application of Cryptology and Information Security (pp. 99-113). Springer Berlin Heidelberg.

Li, H., Shen, F., Shen, C., et al.,: 'Face recognition using linear representation ensembles', Pattern Recognit., 2016, 59, (1), pp. 72–87

Li, S., Jain, A.: 'Handbook of face recognition' (Springer-Verlag London, London, UK, 2011)

Liu, H., Sun, D., Xiong, K., & Qiu, Z. (2011, October). Is fuzzy vault scheme very effective for key binding in biometric cryptosystems? In Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), 2011 International Conference on (pp. 279-284). IEEE.

Loukas, George. Cyber-physical attacks: A growing invisible threat. Butterworth-Heinemann, 2015.

Luh, R., Marschalek, S., Kaiser, M., Janicke, H., & Schrittwieser, S. (2017). Semantics-aware detection of targeted attacks: a survey. Journal of Computer Virology and Hacking Techniques, 13(1), 47-85.

Lykou, Georgia, Dimitrios Moustakas, and Dimitris Gritzalis. "Defending airports from UAS: A survey on cyber-attacks and counter-drone sensing technologies." *Sensors* 20.12 (2020): 3537.

Ma, J., Wang, Q., & Zhao, Z. (2017). SLAE–CPS: Smart lean automation engine enabled by cyber-physical systems technologies. *Sensors*, *17*(7), 1500.

Magureanu, Gabriela, et al. "Towards UML modeling of cyber-physical systems: A case study for gas distribution." *IEEE 8th International Symposium on Intelligent Systems and Informatics*. IEEE, 2010.

Magureanu, Gabriela, Madalin Gavrilescu, and Dan Pescaru. "Validation of static properties in unified modeling language models for cyber physical systems." *Journal of Zhejiang University SCIENCE C* 14.5 (2013): 332-346.

Maier, Mark W. "Architecting principles for systems-of-systems." *Systems Engineering: The Journal of the International Council on Systems Engineering* 1.4 (1998): 267-284.

Mallet, Frédéric, Eugenio Villar, and Fernando Herrera. "MARTE for CPS and CPSoS." *Cyber-Physical System Design from an Architecture Analysis Viewpoint*. Springer, Singapore, 2017. 81-108.

Manshaei, Mohammad Hossein, et al. "Game theory meets network security and privacy." ACM Computing Surveys (CSUR) 45.3 (2013): 1-39.

Mathew, E. (2020). Swarm intelligence for intelligent transport systems: opportunities and challenges. Swarm Intelligence for Resource Management in Internet of Things, 131-145.

Matthews, Brian W. "Comparison of the predicted and observed secondary structure of T4 phage lysozyme." *Biochimica et Biophysica Acta (BBA)-Protein Structure* 405.2 (1975): 442-451.

Martin, A., Doddington, G., Kamm, T., Ordowski, M., & Przybocki, M. (1997). The DET curve in assessment of detection task performance. National Institute of Standards and Technology Gaithersburg MD.

Maxim, Sergievskiy. "N-ary relations of association in class diagrams: design patterns." *International Journal of Advanced Computer Science and Applications* 7.2 (2016): 265-268.

McGraw, R. (2009, August). Risk-adaptable access control (radac). In Privilege (Access) Management Workshop. NIST–National Institute of Standards and Technology–Information Technology Laboratory (Vol. 25, pp. 55-58).

Metz, C. E. (1978, October). Basic principles of ROC analysis. In Seminars in nuclear medicine (Vol. 8, No. 4, pp. 283-298). WB Saunders.

Meyer, Bertrand. "Applying'design by contract'." *Computer* 25.10 (1992): 40-51.

Molloy, Ian, Luke Dickens, Charles Morisset, Pau-Chen Cheng, Jorge Lobo, and Alessandra Russo. "Risk-based security decisions under uncertainty." In Proceedings of the second ACM conference on Data and Application Security and Privacy, pp. 157-168. ACM, 2012

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. "Building state capacity: Evidence from biometric smartcards in India." American Economic Review 106.10 (2016): 2895-2929.

Myerson, Roger B. "Refinements of the Nash equilibrium concept." *International journal of game theory* 7.2 (1978): 73-80.

Nadira, Benlahrache, Chafia Bouanaka, Mohamed Bendjaballah, and Abdoudjallil Djarri. "Towards an UML-based SoS Analysis and Design Process." In *2020 International Conference on Advanced Aspects of Software Engineering (ICAASE)*, pp. 1-8. IEEE, 2020.

Nash, John F*. Non-Cooperative Games*. *Annals of Mathematics* 54.2 (1951).

Nazarenko, Artem A., and Ghazanfar Ali Safdar. "Survey on security and privacy issues in cyber physical systems [J]." *AIMS Electronics and Electrical Engineering* 3.2 (2019): 111-143.

NIST Special Publication 800-63-3, Digital Identity Guidelines, June 2017 (https://doi.org/10.6028/NIST.SP.800-63-3)

National Science Foundation (NSF), Cyber-Physical System (CPS), Program Solicitation NSF 10-515, 2010. Available at http://www.nsf.gov/pubs/2010/nsf10515/nsf10515.htm

Nunamaker, Jay F., et al. "Embodied conversational agent-based kiosk for automated interviewing." Journal of Management Information Systems 28.1 (2011): 17-48

Nuppeney, M., Breitenstein, M., & Niesing, M. (2010). EasyPASS: evaluation of face recognition performance in an operational automated border control system. In Int. Biometric Performance Conf.

OASIS XACML Technical Committee. (2013). eXtensible access control markup language (XACML) Version 3.0. Oasis standard, OASIS. DOI= http://docs.oasisopen. org/xacml/3.0/xacml- 3.0-core-spec-cs-01-en.pdf

Parkhi, O., Vedaldi, A., Zisserman, A.: 'Deep face recognition'. British Machine Vision Conf., Swansea, UK, September 2015

Pedersen, Torben. "Non-interactive and information-theoretic secure verifiable secret sharing." In Advances in Cryptology—CRYPTO'91, pp. 129-140. Springer Berlin/Heidelberg, 1992.

Phillips, P. Jonathon, and Alice J. O'toole. "Comparison of human and computer performance across face recognition experiments." *Image and Vision Computing* 32.1 (2014): 74-85.

Poole, Robert W. Toward risk-based aviation security policy. No. 2008-23. OECD/ITF Joint Transport Research Centre Discussion Paper, 2008.

Raman, Ramakrishnan, and Yogananda Jeppu. "An Approach for formal verification of machine learning based complex systems." *INCOSE International Symposium*. Vol. 29. No. 1. 2019.

Raman, Ramakrishnan, and Yogananda Jeppu. "Does The Complex SoS Have Negative Emergent Behavior? Looking For Violations Formally." *2021 IEEE International Systems Conference (SysCon)*. IEEE, 2021.

Ratha, Nalini K., Jonathan H. Connell, and Ruud M. Bolle. "Enhancing security and privacy in biometrics-based authentication systems." IBM systems Journal 40, no. 3 (2001): 614-634.

Rathgeb, Christian, and Andreas Uhl. "A survey on biometric cryptosystems and cancelable biometrics." EURASIP Journal on Information Security 2011.1 (2011)

Christian Rathgeb, Frank Breitinger, and Christoph Busch. Alignment-free cancelable iris biometric templates based on adaptive bloom filters. In Julian Fierrez, Ajay Kumar, Mayank Vatsa, Raymond N. J.

Veldhuis, and Javier Ortega-Garcia, editors, International Conference on Biometrics – ICB, pages 1–8. IEEE, 2013.

Rathgeb, C., Breitinger, F., Busch, C., & Baier, H. (2014). On application of bloom filters to iris biometrics. IET Biometrics, 3(4), 207-218.

Rescorla, E.: SSL and TLS: Designing and Building Secure Systems. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (2001).

Rivest, R., Shamir, A., Adleman, L.: A method for obtaining digital signatures and public key cryptosystems. Communications of the ACM, 21(2), 120-126 (1978).

Safjanski, T.: Prospects for the Development of the International Criminal Police Organisation Interpol. Internal Security 7.2 (2015).

Sage, Andrew P., and Christopher D. Cuppan. "On the systems engineering and management of systems of systems and federations of systems." Information knowledge systems management 2.4 (2001): 325-345.

Samaria, Ferdinando S., and Andy C. Harter. "Parameterisation of a stochastic model for human face identification." Applications of Computer Vision, 1994. Proceedings of the Second IEEE Workshop on. IEEE, 1994.

Sandhya, M., & Prasad, M. V. (2017). Biometric Template Protection: A Systematic Literature Review of Approaches and Modalities. In Biometric Security and Privacy (pp. 323-370). Springer International Publishing.

Sarier, Nevire Deniz. "Comments on biometric-based non-transferable credentials and their application in blockchain-based identity management." *Computers & Security* 105 (2021): 102243.

Schnorr, Claus-Peter. "Efficient signature generation by smart cards." Journal of cryptology 4.3 (1991): 161-174.

Selic, Bran. "What's New in UML™ 2.0?." (2005).

Selic, Bran. "A Systematic Approach to Domain-Specific Language Design Using UML." *Proceedings of the 10th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing*. 2007.

Sepas-Moghaddam, Alireza, Fernando M. Pereira, and Paulo Lobato Correia. "Face recognition: a novel multi-level taxonomy based survey." *IET Biometrics* 9.2 (2019): 58-67.

Shaikh, Siraj A., and Joseph R. Rabaiotti. "Characteristic trade-offs in designing large-scale biometric-based identity management systems." Journal of Network and Computer Applications 33.3 (2010): 342-351.

Sherali, Hanif D., Allen L. Soyster, and Frederic H. Murphy. "Stackelberg-Nash-Cournot equilibria: characterizations and computations." *Operations Research* 31.2 (1983): 253-276.

Sherif, M.: Protocols for Secure Electronic Commerce.2nd. CRC Press, (2016).

Shi, J., Wan, J., Yan, H., & Suo, H. (2011, November). A survey of cyber-physical systems. In Wireless Communications and Signal Processing (WCSP), 2011 International Conference on (pp. 1-6). IEEE.

Simon, Herbert A. (1947). *Administrative Behavior: a Study of Decision-Making Processes in Administrative Organization* (1st ed.). New York: Macmillan.

Simoens, K., Yang, B., Zhou, X., Beato, F., Busch, C., Newton, E. M., & Preneel, B. (2012, March). Criteria towards metrics for benchmarking template protection algorithms. In Biometrics (ICB), 2012 5th IAPR International Conference on (pp. 498-505). IEEE.

Singh, Pawan. "Aadhaar and data privacy: biometric identification and anxieties of recognition in India." Information, Communication & Society (2019): 1-16.

Sirovich, L., Kirby, M. (1987). Low-dimensional procedure for the characterization of human faces. Josa a, 4(3), 519-524.

Soria Zurita, Nicolás F., and Irem Y. Tumer. "A Survey: Towards Understanding Emergent Behavior in Complex Engineered Systems." *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. Vol. 58219. American Society of Mechanical Engineers, 2017.

Soutar, C., Roberge, D., Stoianov, A., Gilroy, R., & Kumar, B. V. (1998, April). Biometric encryption using image processing. In Photonics West'98 Electronic Imaging (pp. 178-188). International Society for Optics and Photonics.

Sutcu, Y., Li, Q., & Memon, N. (2007). Protecting biometric templates with sketch: Theory and practice. IEEE Transactions on Information Forensics and Security, 2(3), 503-512.

Sutcu, Y., Li, Q., & Memon, N. (2009, May). Design and analysis of fuzzy extractors for faces. In SPIE Defense, Security, and Sensing (pp. 73061X-73061X). International Society for Optics and Photonics.

Srinivas, T. Aditya Sai, Ramasubbareddy Somula, and K. Govinda. "Privacy and security in Aadhaar." Smart Intelligent Computing and Applications. Springer, Singapore, 2020. 405-410.

Talreja, Veeru, Matthew C. Valenti, and Nasser M. Nasrabadi. "Multibiometric secure system based on deep learning." 2017 IEEE Global conference on signal and information processing (globalSIP). IEEE, 2017.

Tan, X., Chen, S., Zhou, Z. H., & Zhang, F. (2006). Face recognition from a single image per person: A survey. Pattern recognition, 39(9), 1725-1745.

A. B. J. Teoh and C. T. Yuang, “Cancellable biometrics realization with multispace random projections,” IEEE Trans. Syst., Man, Cybern. —Special Issue on Recent Advances in Biometrics Systems, vol. 37, no. 5, pp. 1096–1106, Oct. 2007.

Turk, M., & Pentland, A. (1991). Eigenfaces for recognition. Journal of cognitive neuroscience, 3(1), 71-86.

Turk, M. A., & Pentland, A. P. (1991, June). Face recognition using eigenfaces. In Computer Vision and Pattern Recognition, 1991. Proceedings CVPR'91., IEEE Computer Society Conference on (pp. 586-591). IEEE.

P. Tuyls and J. Goseling, Capacity and Examples of Template Protecting, Biometric Authentication Systems, Biometric Authentication Workshop (BioAW 2004), LNCS 3087, 158–170, Prague, 2004.

Uludag U, Pankanti S, Prabhakar S, Jain AK: Biometric cryptosystems: issues and challenges. Proc IEEE 2004, 92(6):948-960.

Van Laarhoven, Peter JM, and Emile HL Aarts. "Simulated annealing." *Simulated annealing: Theory and applications*. Springer, Dordrecht, 1987. 7-15.

Verheul, E. R. (2001, December). Self-blindable credential certificates from the Weil pairing. In International Conference on the Theory and Application of Cryptology and Information Security (pp. 533-551). Springer Berlin Heidelberg.

Veseli, Fatbardh, and Jetzabel Serna. "Evaluation of privacy-ABC technologies-a study on the computational efficiency." IFIP International Conference on Trust Management. Springer, Cham, 2016.

Viola, Paul, and Michael Jones. "Fast and robust classification using asymmetric adaboost and a detector cascade." *Advances in Neural Information Processing System* 14 (2001).

Vu, An Hoa, et al. "Cybersage: A tool for automatic security assessment of cyber-physical systems." *International Conference on Quantitative Evaluation of Systems*. Springer, Cham, 2014.

F.-Y. Wang, “The emergence of intelligent enterprises: From CPS to CPSS,” IEEE Intell. Syst., vol. 25, no. 4, pp. 85–88, Jul./Aug. 2010.

Wang, Fei-Yue, et al. "Parallel driving in CPSS: A unified approach for transport automation and vehicle intelligence." *IEEE/CAA Journal of Automatica Sinica* 4.4 (2017): 577-587.

Wang, E. K., Ye, Y., Xu, X., Yiu, S. M., Hui, L. C. K., & Chow, K. P. (2010, December). Security issues and challenges for cyber physical system. In Proceedings of the 2010 IEEE/ACM Int'l Conference on Green Computing and Communications & Int'l Conference on Cyber, Physical and Social Computing (pp. 733-738). IEEE Computer Society.

Wang, Y., & Plataniotis, K. N. (2010). An analysis of random projection for changeable and privacy-preserving biometric verification. IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), 40(5), 1280-1293.

World Customs Organization, International Air Travel Association, International Civil Aviation Organization, GUIDELINES ON ADVANCE PASSENGER INFORMATION (API), http://www.un.org/sc/ctc/pdf/APIGuidelines.pdf 2013

World Economic Forum, “Known Traveller Digital Identity” https://ktdi.org/ (accessed Nov. 2021)

Wilson, K. (2016). Gone With the Wind?: The Inherent Conflict between API/PNR and Privacy Rights in an Increasingly Security-Conscious World. Air and Space Law, 41(3), 229-264.

Wirfs-Brock, Rebecca J., and Ralph E. Johnson. "Surveying current research in object-oriented design." *Communications of the ACM* 33.9 (1990): 104-124.

Wiskott, L., Fellous, J., Kruger, N., et al.,: ‘Face recognition by elastic bunch graph matching’. Int. Conf. on Image Processing, Santa Barbara, CA, USA, October 1997

Wu, Z., Wang, Y., Pan, G.: ‘3D face recognition using local shape map’. Int. Conf. on Image Processing, Singapore, October 2004

Xiong, Gang, et al. "Cyber-physical-social system in intelligent transportation." IEEE/CAA Journal of Automatica Sinica 2.3 (2015): 320-333.

ODroid XU4 product page http://www.hardkernel.com/main/products/prdt_info.php, last accessed 2018/09/12.

Yourdon, Edward. *Modern structured analysis*. Yourdon press, 1989.

Yu, Chung-Hsien, et al. "Crime forecasting using data mining techniques." *2011 IEEE 11th international conference on data mining workshops*. IEEE, 2011.

Zeng, Jing, et al. "A survey: Cyber-physical-social systems and their system-level design methodology." Future Generation Computer Systems 105 (2020): 1028-1042.

Zhang, T., Wang, B., Li, F., et al.,: ‘Decision pyramid classifier for face recognition under complex variations using single sample per person’, Pattern Recognit., 2017, 64, (1), pp. 305–313

Zhao, W., Chellappa, R., Phillips, P. J., & Rosenfeld, A. (2003). Face recognition: A literature survey. ACM computing surveys (CSUR), 35(4), 399-458.

[Appendix 1] Configuration and Summaries

The following Configuration corresponds to

```
transitions2 = {\
                "S_mtc_begin": {"n1_mtc_cloud": 1},
                "n1_mtc_cloud": {"n22_dep_manual": 0.9,  "e1_mtc_denied": 0.1},
                "e1_mtc_denied": {},
                "n21_dep_kiosk": {"n22_dep_manual": 0.2, "e2_no_board": 0.1, "n3_in_transit":
0.7},
                "n22_dep_manual":{"e2_no_board": 0.2, "n3_in_transit": 0.8,},
                "e2_no_board":{},
                "n3_in_transit": {"e3_disruption": 0.003, "n41_arr_kiosk": 0.967,
"n42_arr_secondary": 0.03},
                "e4_arr_deported":{},
                "e3_disruption":{},
                "n41_arr_kiosk":{"E_arr_admitted": 0.7,"n42_arr_secondary": 0.3},
                "n42_arr_secondary":{"e4_arr_deported": 0.25, "E_arr_admitted" : 0.75},
                "E_arr_admitted":{}}
```

Summary

```
Counter({'E_arr_admitted': 665, 'e2_no_board': 169, 'e1_mtc_denied': 110, 'e4_arr_deported': 55,
'e3_disruption': 1})
49160
E_arr_admitted
56970
e4_arr_deported
-1310
e1_mtc_denied
-1540
e2_no_board
-4901
e3_disruption
-59
{'E_arr_admitted': 56970, 'e4_arr_deported': -1310, 'e1_mtc_denied': -1540, 'e2_no_board': -4901,
'e3_disruption': -59}
```

[Appendix 2] Source Code

```
import pandas as pd
import numpy as np
from collections import Counter

# 1) Data Definitons
# transitions and costs variables map in value and in structure to tables   in the thesis

scX = "Sc1"


transitions = { \
"t01":( 1.00, 1.00, 1.00),
"t11":( 0.75, 0.90, 0.90),
"t12":( 0.25, 0.10, 0.10),
"t21":( 0.60, 0.60, 0.80),
"t22":( 0.40, 0.40, 0.20),
"t23":( 0.80, 0.80, 0.90),
"t24":( 0.20, 0.20, 0.10),
"t31":( 0.91, 0.91, 0.999),
"t32":( 0.09, 0.09, 0.001),
"t41":( 0.85, 0.95, 0.95),
"t42":( 0.15, 0.05, 0.05),
"t43":( 0.75, 0.40, 0.40),
"t44":( 0.25, 0.60, 0.60),
"t51":( 0.95, 0.96, 0.99),
"t52":( 0.05, 0.04, 0.01)}
costs = {\
"S_start":(  -2,  -2,  -2),
"p11_TA issue":( -20,  -8,  -8),
"e1_TA denied" :( -15, -16, -16),
"p21_preboard_main":(  -9,  -9,  -3),
"p22_preboard_overflow":( -10, -10,  -4),
"e2_no_board" :( -20, -20, -20),
"p31_inflight":(  50,  50,  50),
"e3_inflight_disrupt" :(-200,-200,-200),
"p41_arrival_primary":( -15,  -5,  -5),
"p42_arrival secondary":( -15, -15, -15),
"e4_arrival_deported" :( -75, -75, -75),
"p5_admitted" :(  80,  80,  80),
"E1_compliant_stay" :( 200, 200, 200),
"E2_disruptive_stay" :(-500,-500,-500)}

# 2) Adaptors
# These definitions and methods adapt from the thesis representation to the data structures
required by the simulation
# make_concourse() is the method that does the work
transition_defs = {\
"t01" :(    "S_start" ,"p11_TA issue"),
"t11" :( "p11_TA issue","p21_preboard_main"),
"t12" :( "p11_TA issue", "e1_TA denied"),
"t21" :( "p21_preboard_main" ,"p31_inflight"),
"t22" :( "p21_preboard_main" ,"p22_preboard_overflow"),
"t23" :( "p22_preboard_overflow" ,"p31_inflight"),
"t24" :( "p22_preboard_overflow" , "e2_no_board"),
"t31" :( "p31_inflight" ,"p41_arrival_primary"),
"t32" :( "p31_inflight", "e3_inflight_disrupt"),
"t41" :( "p41_arrival_primary", "p5_admitted"),
"t42" :( "p41_arrival_primary","p42_arrival secondary"),
"t43" :( "p42_arrival secondary", "p5_admitted"),
"t44" :( "p42_arrival secondary", "e4_arrival_deported"),
"t51" :( "p5_admitted", "E1_compliant_stay"),
```

```
"t52" :( "p5_admitted", "E2_disruptive_stay"),}



#I know python has a logger. This is easier.
class LogFile:
    filename = None
    file = None
    @classmethod
    def open(self, fname):
        self.filename = fname
        self.file = open(fname, 'w')
    @classmethod
    def write(self, string):
        self.file.write( string )
#        print(string)

    @classmethod
    def writelines(self, string):
        self.file.writelines( string )
#        print(string)

    @classmethod
    def close(self):
        if not self.file is None:
            self.file.close()
        self.file = None

    @classmethod
    def get_name(self):
        return self.filename
def make_concourse(sc_id_str = "Sc3"):
indexes = {"Sc1":0, "Sc2": 1, "Sc3": 2  }
index = indexes[sc_id_str]
concourse = {start : { } for tr, (start, end) in transition_defs.items()}
scX_transitions = {key : value[index] for key, value in transitions.items()}
scX_costs = {key : value[index] for key, value in costs.items()}
for transition_id, (start_node, end_node) in transition_defs.items():
try:
concourse[end_node]
except KeyError:
concourse[end_node] = {}
concourse[start_node][end_node] = scX_transitions[transition_id]
return (scX_costs, transition_defs, scX_transitions, concourse )

# 3) Run loop
# This method implements the driver harness to runs the selected scenario




def run_all_scenarios(loop = 1000 ):
    scenarios = ("Sc1", "Sc2", "Sc3")
    exits = ("e1_TA denied", "e2_no_board", "e3_inflight_disrupt", "e4_arrival_deported",
"E1_compliant_stay", "E2_disruptive_stay")
    exit_headers = ("e1", "e2", "e3", "e4", "E1", "E2")
    results = {}



    for scX in scenarios:
        ( f, c) = run_experiment(loop, scX = scX )
        results[scX] = ( f, c)

    print( "************Comparative Scenario Summary************\n" , end = "" )
    for exit in exit_headers:
        print( "\t" + exit , end = "")

    for scX in scenarios:
        print( "\n" + scX , end = "")
```

```
        total_travelers = 0
        total_utility = 0

        for exit in exits:
            value = results[scX][0].get(exit,0)
            total_travelers += value
            print( "\t" + str(value), end = "")
        print( "\t" + str( total_travelers ), end = "")
        print( "\n",  end = "")
        for exit in exits:
            value = results[scX][1].get(exit,0)
            total_utility += value
            print( "\t" + str(value ), end = "" )
        print( "\t" + str( total_utility ), end = "" )

    print("\n==================LEGEND=====================")
    print("\n------------1) Exit Abbreviations:------------")
    for i in range(0, len(exit_headers)):
        print( exit_headers[i].ljust(5) +  ":" + exits[i] )

    print("\n------------2)Comparative Costs:------------")
    for (k,v) in costs.items():
        print(k.ljust(20) + "\t:" + str ( v ))


    print("\n------------3)Comparative Transitions:------------")
    for (k,v) in transitions.items():
        print(k.ljust(5) + ":" + str ( v ))



    print("\n------------4)Transition Definitions:------------")
    for (k,v) in transition_defs.items():
        print(k.ljust(5) + ":" + str ( v ))

    print( "*******************************************"  )



def run_experiment( loop = 10, scX = "Sc2", seed = 2021):
    filename = "thesis_simulation_out.txt"
    (costs, debug_tr_defs1, debug_tr_defs2, transitions ) =  make_concourse( scX )
    filename = scX + "_" + filename
    LogFile.open(filename)
    print(costs)
    print(debug_tr_defs1)
    print(debug_tr_defs2)
    print(transitions )
    print("---------------")

    LogFile.write( scX +"\n")
    LogFile.write( "n=" + str(loop) +"\n")

    LogFile.write( str(costs) +"\n")
    LogFile.write( str(transitions) +"\n")
    LogFile.write("subject_id" + "\t" + "state" + "\t" + "cost" +"\n")
    r =run_loop(start = "S_start", stt = transitions,  entry_costs = costs, i = loop)
    s, c = summarize_all(r)
    print("//////////////////////")
    print(c)
    print("\\\\\\\\\\\\\\\\\\\\\\\\")
    (fr, cs) = plot(s)
    LogFile.close()
    print(filename)
    return (fr, cs)



def run_loop(start, stt , entry_costs, i = 10, seed = 2021):
    rst = np.random.RandomState(seed)
    result = {}
```

```
    for j in rst.randint(1000000, 100000001, i):
        result[j] = traverse_stt(rst = rst, stt2 = stt, costs = entry_costs, subject_id = j,
start = start )
    return result

def traverse_stt(rst, stt2, costs, subject_id, start ):

    result = pd.DataFrame(columns = ["subject_id", "state","cost"])
    current_state = start

    while True:
        new_entry = {"subject_id":subject_id, "state":current_state, "cost":costs[current_state]}
        result = result.append( new_entry, ignore_index=True)
        LogFile.writelines(str(subject_id) + "\t" + str(current_state)  + "\t" +
str(costs[current_state]) +"\n")
        transition_dict = stt2[current_state]

        if( len(transition_dict) == 0):
            break
        else:
            states_names = list(transition_dict.keys())
            transition_probabilities =   list(transition_dict.values())
            current_state = rst.choice(states_names, p = transition_probabilities)
    return result


def summarize_trip(trip_result, result = None):
    subject_id = list(trip_result["subject_id"])[-1]
    last_state = list(trip_result["state"])[-1]
    cost = sum(list(trip_result["cost"]))
#    f.writelines( trip_result )
    print(trip_result)
    new_entry = {"subject_id":subject_id, "last_state":last_state,"total_cost":cost}
    LogFile.writelines( str(new_entry ) + "\n" )

#    data = [subject_id, last_state, cost ]
    if result is None:
        result = pd.DataFrame(columns = ["subject_id", "last_state","total_cost"])
    result = result.append( new_entry, ignore_index=True)
    return result


def summarize_all(loop_result):
    result = None
    counter = Counter()
    for trip_df in list(loop_result.values()):
        result = summarize_trip(trip_df, result)
        counter.update(trip_df["state"])

    return ( result, counter )

def plot(summary):
    last_state_bag = summary["last_state"]
    frequencies = Counter(last_state_bag)
    df = pd.DataFrame.from_dict(frequencies, orient='index')
    df.plot(kind = "bar")

    print(frequencies)
    print("Total cost:" + str(sum(summary["total_cost"])))
    costs_by_end_state = cost_by_end_state(summary)
    print()
    LogFile.writelines( str( frequencies ) +"\n" )

    LogFile.write( "Total cost:" + str(sum(summary["total_cost"] )) +"\n" )
    LogFile.writelines( str(cost_by_end_state(summary)) +"\n" )
    return (frequencies, costs_by_end_state )

def cost_by_end_state(summary):
    last_state_bag = summary["last_state"]
    freqencies = Counter(last_state_bag)
    cost_summary = {}
```

```
    for k in list(freqencies.keys()):
        print(k)
        LogFile.writelines( "cost_by_end_state" +"\t"  )
#       f.writelines( k +"\n" )
        df_all_for_last_state_k = summary[summary["last_state"] == k ]
        #print(df_all_k)
        s = sum(df_all_for_last_state_k["total_cost"])
        print(s)
        LogFile.write( str(s) +"\n" )
        cost_summary[k] = s
    return cost_summary
```

[Appendix 3] Sample Data

The following is a except of sample data from a simulation execution:

```
53520052 S_mtc_begin -2
53520052 n1_mtc_cloud -2
53520052 n22_dep_manual -5
53520052 n3_in_transit 50
53520052 n41_arr_kiosk -2
53520052 E_arr_admitted 50
76722069 S_mtc_begin -2
76722069 n1_mtc_cloud -2
76722069 n22_dep_manual -5
76722069 n3_in_transit 50
76722069 n41_arr_kiosk -2
76722069 E_arr_admitted 50
63789689 S_mtc_begin -2
63789689 n1_mtc_cloud -2
63789689 n22_dep_manual -5
63789689 n3_in_transit 50
63789689 n41_arr_kiosk -2
63789689 E_arr_admitted 50
29787840 S_mtc_begin -2
29787840 n1_mtc_cloud -2
29787840 n22_dep_manual -5
29787840 n3_in_transit 50
29787840 n41_arr_kiosk -2
29787840 n42_arr_secondary -13
29787840 e4_arr_deported -50
60902573 S_mtc_begin -2
60902573 n1_mtc_cloud -2
60902573 n22_dep_manual -5
60902573 n3_in_transit 50
60902573 n41_arr_kiosk -2
60902573 E_arr_admitted 50
3772830 S_mtc_begin -2
3772830 n1_mtc_cloud -2
3772830 n22_dep_manual -5
3772830 n3_in_transit 50
3772830 n41_arr_kiosk -2
3772830 E_arr_admitted 50
1743382 S_mtc_begin -2
1743382 n1_mtc_cloud -2
1743382 n22_dep_manual -5
1743382 n3_in_transit 50
1743382 n41_arr_kiosk -2
1743382 E_arr_admitted 50
67484268 S_mtc_begin -2
67484268 n1_mtc_cloud -2
67484268 n22_dep_manual -5
67484268 n3_in_transit 50
67484268 n41_arr_kiosk -2
67484268 n42_arr_secondary -13
67484268 E_arr_admitted 50
69772670 S_mtc_begin -2
69772670 n1_mtc_cloud -2
69772670 e1_mtc_denied -10
14582299 S_mtc_begin -2
14582299 n1_mtc_cloud -2
14582299 e1_mtc_denied -10
1283293 S_mtc_begin -2
1283293 n1_mtc_cloud -2
1283293 n22_dep_manual -5
1283293 n3_in_transit 50
1283293 n41_arr_kiosk -2
```

```
1283293 E_arr_admitted 50
98680917 S_mtc_begin -2
98680917 n1_mtc_cloud -2
98680917 n22_dep_manual -5
98680917 n3_in_transit 50
98680917 n41_arr_kiosk -2
98680917 E_arr_admitted 50
36637597 S_mtc_begin -2
36637597 n1_mtc_cloud -2
36637597 n22_dep_manual -5
36637597 n3_in_transit 50
36637597 n41_arr_kiosk -2
36637597 E_arr_admitted 50
13564632 S_mtc_begin -2
13564632 n1_mtc_cloud -2
13564632 e1_mtc_denied -10
26385462 S_mtc_begin -2
26385462 n1_mtc_cloud -2
26385462 n22_dep_manual -5
26385462 n3_in_transit 50
26385462 n41_arr_kiosk -2
26385462 n42_arr_secondary -13
26385462 E_arr_admitted 50
89850630 S_mtc_begin -2
89850630 n1_mtc_cloud -2
89850630 n22_dep_manual -5
89850630 e2_no_board -20
14143142 S_mtc_begin -2
14143142 n1_mtc_cloud -2
14143142 n22_dep_manual -5
14143142 n3_in_transit 50
14143142 n41_arr_kiosk -2
14143142 n42_arr_secondary -13
14143142 E_arr_admitted 50
63754886 S_mtc_begin -2
63754886 n1_mtc_cloud -2
63754886 n22_dep_manual -5
63754886 n3_in_transit 50
63754886 n41_arr_kiosk -2
63754886 n42_arr_secondary -13
63754886 E_arr_admitted 50
14505057 S_mtc_begin -2
14505057 n1_mtc_cloud -2
14505057 n22_dep_manual -5
14505057 e2_no_board -20
89980069 S_mtc_begin -2
89980069 n1_mtc_cloud -2
89980069 n22_dep_manual -5
89980069 n3_in_transit 50
89980069 n41_arr_kiosk -2
89980069 E_arr_admitted 50
98957244 S_mtc_begin -2
98957244 n1_mtc_cloud -2
98957244 n22_dep_manual -5
98957244 n3_in_transit 50
98957244 n41_arr_kiosk -2
98957244 E_arr_admitted 50
85899557 S_mtc_begin -2
85899557 n1_mtc_cloud -2
85899557 n22_dep_manual -5
85899557 n3_in_transit 50
85899557 n41_arr_kiosk -2
85899557 n42_arr_secondary -13
85899557 E_arr_admitted 50
56869122 S_mtc_begin -2
56869122 n1_mtc_cloud -2
56869122 n22_dep_manual -5
56869122 n3_in_transit 50
56869122 n41_arr_kiosk -2
56869122 E_arr_admitted 5
```