# ERC Consolidator Grant 2016
# Research proposal [Part B2]

## Part B2: The Scientific Proposal

*Article 7 :* ***Respect for private and family life***
*Everyone has the right to respect for his or her private and family life, home and communications.*
*Article 8 :* ***Protection of personal data***
*1. Everyone has the right to the protection of personal data concerning him or her.*
CHARTER OF FUNDAMENTAL RIGHTS OF THE EUROPEAN UNION [1]

## Section a: State-of-the-art and objectives

Note: Section numbers continue from part B1 to allow for consistent cross-referencing and reduce redundancy.

### 6 Introduction

The **primary objective** of Digidow is to **enable security-relevant interactions in the physical world without carrying any physical objects to certify the individual's identity**. This major step depends on secure biometric authentication to use the individual's body, behavior, and habits for authentication instead of physical tokens such as passports. Fulfilling this vision will enable further breakthroughs in ubiquitous/pervasive, mobile, wearable, and embedded computing, human/computer interaction, or more specific application areas from other disciplines such as social or legal sciences.

This goal of object-less authentication in the physical world demands a counterpart in the digital world that supports individuals' interactions with the required identification and authentication data as well as state information — we call this instance the *Digital Shadow*. Although biometric authentication and identity data management could be more easily implemented with centralized databases, the associated danger of abuse is too high to be acceptable. The **secondary objective** is therefore to **support secure authentication with decentralized personal agents**, which remain under full control of each associated individual and manage their digital identities and assets. It is clear that trustworthy digital identity documents (sometimes called electronic ID or e-ID) such as passports will, during their *creation* process, still be signed by central organizational bodies such as countries in a similar way as their current physical manifestations are created now. However, the major difference in a decentralized architecture is that during their *usage*, control remains with the individual and not with the organization hosting the centralized databases.

In the sense of managing different forms of digital identity, personal agents in Digidow loosely fulfill the role of current password managers, although with a significantly wider scope and a completely different user interaction model. The general area of authentication options for and use of digital identity for purely digital services are currently being investigated in different research and development projects. This includes the EU projects STORK [2] for using digital services with e-ID across Europe (targeting mostly government use) and FutureID [3] for e-ID desktop use cases (targeting free market use) as well as national initiatives like the Austrian citizenship card (Bürgerkarte), Estonian e-Residency, Finnish identity card, or German ID card (Neuer Personalausweis, nPA). All these projects are orthogonal to Digidow in their use of digital identity for digital services; however, focusing on services in the physical world and specifically without using physical smart cards for authentication requires completely novel approaches to secure code and infrastructure. After successful authentication to a personal agent in Digidow – which could act like a physical smart card in terms of cryptographic protocol exchange – the results of STORK, FutureID and others can directly build upon the authentication state for using digital services. The recently concluded projects NewP@ss [4] and FastPass [5] for Automated Border Control gates at European borders come closest in their use of the ICAO passport standard and use case of crossing borders, but again rely on physical objects (passports) with embedded RFID storage.

In the following, we structure the discussion of the state of the art by the four main areas to be solved.

## 7 Area A: Authentication

In terms of biometric authentication, the most compelling modalities for Digidow currently seem to be face, fingerprint, and iris recognition. We will build upon the current state-of-the-art in terms of hardware sensors, feature extraction, and classification / error metrics, and thus in the following review only the most recent and relevant results for their applicability to our use cases.

No single approach can currently be considered the most powerful state-of-the-art technology in face recognition and face authentication [6]. Current face recognition approaches are combinations of a huge diversity of different face recognition concepts, based on both geometric facial features and appearance based face recognition [7]. The underlying mechanisms for geometric feature based face recognition include concepts from morphable models [8] to processing features using Hidden Markov Models (HMM) [9]. Appearance based face recognition approaches include concepts from e.g. PCA/Eigenface and LDA/Fisherface or Gabor Wavelets [10, 11] for (pre)processing and apply diverse machine learning models, such as neural networks (NN), support vector machines (SVM), boosting, or combinations [12, 13, 14]. Error rates reported in individual publications strongly depend on the dataset chosen for evaluation and the evaluation setup [6, 15]. Including different types of data, such as range data to obtain 2+1D/3D data, in combination with optical (color) information can further improve results [16, 17]. Also, the recently emerged and heavily discussed topic of deep learning seems promising for face recognition applications [18, 19] as it already outperformed other machine learning models in related recognition tasks [20]. Face recognition is a highly important field for Digidow, as we will use it both as a (weak) biometric authenticator and for tracking individuals using camera networks.

Fingerprint recognition is a mature topic in the sense that different kinds of sensors, features, and matching metrics have been analyzed and accurate settings are known and widely available in off-the-shelf products. The same is mostly true for iris recognition, although the optical sensors are not fully practical in mobile settings.

We may work on specific problems such as face recognition under poor lighting conditions or with seasonal diversity (hats, beards, tanned vs. pale skin, etc.) by utilizing deep learning approaches. The complementary issue of multi-party authentication in the sense of biometric template storage (in the personal agent), sensor, and verifier being under three separate realms of administrative control has not been tackled before and will be our main focus (see section 13). This will require a **fundamentally novel separation of the standard data analysis pipeline** in biometric authentication systems with **trust spread over multiple systems**.

## 8 Area B: Tracking

The physical location of a person at a certain point in time is considered highly sensitive information. Throughout their everyday lives, people are moving in a (mostly) unique pattern and leave a fingerprint of visited locations that might be used to conclude on their identity (cf. [21, 22, 23]). Finding (and most often recording) these traces is commonly referred to as tracking. Systems to track people outdoors are widespread in modern societies, typically based on Global Navigation Satellite Systems that allow accuracies of below one meter of positioning error (using e.g. DGPS or RTK techniques [24]). Indoor tracking demands a more complex approach in terms of positioning and tracking algorithms. For such scenarios, systems have been developed that make use of various different sensor inputs, including radio signal-based sensors (e.g. UWB [25], Pseudolites [26], WiFi [27, 28], RFID [29] and Bluetooth LE [30]), computer vision approaches (using single cameras [31] or visual sensor networks [32]), inertial measurements (e.g. pedestrian dead reckoning [33]), acoustical sensors [34], or a combination of some of these technologies often referred to as sensor fusion (e.g. [35]). As their outdoor counterparts, these systems also achieve accuracy results of one meter or better.

From a hardware point of view, tracking systems can be separated into two groups: client-based systems that require their users to carry a certain mobile device for being localized, and infrastructure-based systems that make use of tracking sensors deployed in the environment (cf. [36]). With the coming of age of mobile computing, client-based tracking systems gained popularity, not least because – in theory (cf. [23]) – it is easier for such systems to ensure privacy protection. For Digidow, we explicitly aim at device-free scenarios and therefore aim for comparable privacy guarantees in infrastructure-based systems.

In [37], Youssef et al. coined the term Device-free Passive (DfP) localization systems describing tracking solutions that explicitly exclude user-worn sensors. For realizing such DfP systems, two types of technologies are currently being investigated. The first technology is again based on radio signals. Several systems have

been presented that scan the radio frequency (RF) field for changes to conclude on people's movements, e.g. Ichnaea [38], Nuzzer [39], tracking pedestrian flows within buildings on airports [40], or determining the activities of people present in the RF field [41]. The open problem is differentiating between people.

The second technology is based on multi-camera sensor systems, often referred to as visual sensor networks (VSN). Due to efficient face recognition algorithms for video streams (e.g. [42]), such systems have the potential to automatically identify individuals. Once identified, tracking a person is realized through collaborating nodes within the VSN that use wireless communication for continuous surveillance (cf. [43, 44]). The FireFly Mosaic system [45] is an example of collaborating camera nodes used to track the activities of a person in their home environment. In [46], an algorithm for tracking multiple persons in a complex environment using synchronized video streams has been presented, which was claimed to achieve metrically accurate position estimates. Recent projects in this field of research dealt with the problem of covering a wide area of interest with only sparse camera node coverage. In [47], Song et al. summarized algorithms and methods for visual surveillance systems covering the topics of intra-camera tracking, inter-camera tracking, camera relationships, and global activity understanding. Given the potential of collaborating VSN nodes and RF-based tracking infrastructure for automated, efficient and large-scale tracking, people might be tracked wherever they go [48].

In Digidow, we explicitly focus on **privacy-conscious tracking**. In terms of VSN applications, the TrustEYE.M4 sensor nodes [49] address some issues. The authors present the concept of sensor-level privacy protection that intends to bring security and privacy down to the lowest level of sensor systems in order to reduce the number of other components that need to be trusted. The AnonySense framework [50] also intends to preserve privacy in ubiquitous tracking scenarios by making use of intermediate services to obfuscate user-related data. We will adopt the basic concept of privacy protection in sensors (not in third-party infrastructure), but need to fundamentally extend it to multiple parties in a 3-way communication.

## 9 Area C: Trust

The vision of a device-free, decentralized identity system requires to establish a trust relationship between personal agents, biometric sensors and the verifiers as well as between individuals and biometric sensors. For the latter, it is required to provide the user a continuous feedback of the trustworthiness of the used sensor, which is an open issue when considering the lack of trusted user interfaces.

Establishing trust between two computing devices (i.e. personal agents, biometric sensors, and verifiers — but not individuals) often includes the usage of additional tamper resistant hardware as a root of trust. Examples for such supplementary hardware are smart cards, secure elements, or trusted platform modules (TPM). They consist of an auxiliary processor as well as tamper resistant storage to store and process security critical data and keys directly on the chip. Use cases for this hardware range from mobile device scenarios (e.g. our proposed ecosystem for mobile devices [51]), credential storage for remote services (e.g. keys for company VPN access), or different kinds of identification system (e.g. the anonymous identity system by Bichsel et al. [52]).

The state-of-the-art technology to establish trust in personal computers (PC) is the usage of a TPM module, which has been specified by the Trusted Computing Group (TCG) [53]. TPMs record the state of the systems with multiple so-called *Platform Configuration Registers (PCR)*, i.e. cumulative hash measurements of software running on the host machine. The whole process of extending the PCR value is usually done concurrently with the boot process of the PC and can be used for data protection (i.e. encrypting/decrypting data only if system runs in a specific state) or for reporting the system state to a remote service with two possible techniques: *Remote Attestation* and *Direct Anonymous Attestation*. We further enhanced the concept of reporting the system state in [54] by proposing an extended and practical approach for mobile devices, which also gives more flexibility to users. Our so-called *certified boot* merges the TPM attestation technique with a signature-based software verification using tamper resistant hardware available in mobile devices.

While the TPM has already been widely deployed on personal computers, there is currently also research on using this technology in other computing devices. The previously mentioned TrustEYE.M4 sensor [49] as well as the TrustCAM [55] use a TPM to verify the integrity, authenticity, and the timestamp of camera pictures to applications or remote services. They propose a concept that establishes trust in sensors at the lowest layers of a computing or sensor system. Similarly, Saroiu and Wolman [56] as well as Gilbert et al. [57] present trustworthy mobile sensing platforms to assure the authenticity of sensor data by using a TPM as trust-anchor.

Nyman et al. [58] also provide an enhanced concept of using the TPM in a system requiring a high level of security and trust to provide the techniques for an electronic identity (e-ID) architecture. This concept still requires the user to carry around a mobile device and always have it functional in case of identity verification. In Digidow, we will significantly extend these techniques for the **paradigm change of device-free usage**.

The *Fast IDentity Online (FIDO) Alliance* is a non-profit industry consortium which develops specifications for user authentication with less reliance on passwords. In their publicly available specification, they define two main protocols: UAF and U2F. The *Universal Authentication Framework (UAF)* [59] protocol can be used by online services for a password-less and multi-factor login while using local authentication mechanisms. The Universal 2nd Factor (U2F) [60] protocol describes the extension of existing authentication infrastructure with a second factor (such as a hardware token). U2F could provide a basis for the identity federation protocol in step 7 (cf. Figure 1 in part B1), but will need to be merged with DAA and extended for 3-party authentication.

## 10 Area D: Network

Most research work on location-based routing and service discovery has been done in the context of mobile ad-hoc networks (MANET) or robot networks. Neither is applicable to Digidow, as we do not assume personal agents to be mobile (physically or digitally), but that they are hosted statically and only communicate with a world-wide network of biometric sensors and verifiers, both of which can be modeled as (network) services. Hence, the service functionality relevant for Digidow is either provided by the personal agent (e.g. authentication, granting of payments, etc.) or a random verifying entity using sensors (e.g. verification of identity, location, etc.). No matter which is hosting the target service, the other party is interested in efficiently finding its counterpart for the event or action that should take place. In this context, the personal agent and the verifying entity form a loosely-coupled multiagent system (cf. [61]), having each peer acting autonomously and proactively to achieve its goals, e.g. tracking the individual (the duty of the personal agent), or authorizing a transaction (the task of the infrastructure).

Discovering the target service is a matter of globally disseminating a) meta data describing the service paired with b) routing information that is considered as slowly changing over time due to the above stated restrictions. In the context of peer-to-peer (P2P) systems this issue is well understood (cf. [62], [63]). The GloServ architecture [64] is one example for realizing global service discovery over P2P networks. For Digidow however, we aim for a system that is not restricted to P2P networking paradigms.

Following the basic idea of the Domain Name System (DNS), we propose a decentralized and privacy-preserving architecture that exploits the physically static nature of the digital entities actively contributing to Digidow. In [65], a hybrid naming system for tracing objects on a global scale has been discussed that used distributed hash tables (DHTs) and DNS for such purposes. The ERGOT system [66] was also based on DHTs to represent services on top of a semantic overlay network (SON). The privacy aspect has not been considered in both approaches. In the context of the digital currency Bitcoin [67], several studies evaluated the impact of globally interacting, decentralized services (e.g. [68]) on the underlying network. Zerocoin [69] demonstrates how to mitigate the privacy issues still present in the original Bitcoin concept.

Briefly summarizing, a massive body of work exists on manifold topics of networking protocols in general and for service discovery in particular. Global service discovery is solved by DNS, but this is neither fully decentralized nor private. For globally, decentralized communication, we can learn from various DHT-based networks (including Kademlia-based "Mainline-DHT" for Bittorrent [70]) and other deployed P2P networks (e.g. Bitcoin [67]). For global networks with privacy guarantees, Tor [71] is currently the best example, while DHTs do not typically provide anonymity [72]. Although initial work on combining flat DHT with hierarchical DNS indexing has been done, bottlenecks still exist [65]. An additional requirement is geospatial service discovery to connect physical locations of individuals with digital services they may want to interact with (a summary can be found in [73]). Our challenge for Digidow is to develop a **service discovery architecture** that fulfills all the requirements (**global, decentralized, low latency, and with privacy guarantees**).

## 11 Summary of challenges in the context of state-of-the-art

In order to reach our main objectives (**secure biometric authentication without carrying physical objects and with decentralized agents under control of each individual**), we identify the following major open issues that are complementary to and not addressed by current state-of-the-art in the four different areas. These

are all associated to the two fundamental challenges discussed before (*scale* and *trust*, cf. section 3 in part B1):

**Multi-party biometric authentication** (*trust*) requires re-defining a data analysis pipeline split over multiple systems with different realms of administrative control and therefore across trust boundaries, and tightly coupling it with a cryptographic authentication protocol.

**Trust in personal agents** (*trust*) requires highly novel results on secure code (of the agent itself), its execution environment (e.g. a cloud provider), and DAA protocol approaches when applied to virtual machine code. Solving the issue of securely running personal agents on remote data centers will also have a huge impact on secure cloud services in general. Although we will focus on designed-to-be-simple code, the developed techniques will in large parts be applicable to arbitrary cloud applications.

**Trust in biometric sensors** (*trust, scale*) requires, based on current state-of-the-art, adding a TPM to each sensor. This does not seem realistic with currently deployed infrastructure, and we therefore additionally require a method to integrate legacy (only partially trusted) sensors into the architecture.

**Trust in verifiers** (*trust, scale*) is not realistic considering an open-ended set of parties acting on an individual's identity. We therefore require novel protocols and privacy-preserving pseudonym derivation methods to deal with the assumption of colluding verifiers that try to monitor and track a global population.

**Service discovery** (*scale*) with latencies in the range of few seconds remains difficult on a global level, and our current best approach is a completely novel combination of geographic filtering with individual user tracking (cf. section 12).

**Backup** (*trust, scale*) is difficult both for the data held by personal agents (how to deal with a temporary or permanent unavailability of the current instance of the agent) considering liveness and authenticity requirements and for authentication methods (as a user, how to securely authenticate to your own remote personal agent when biometric authentication is currently impossible). Current state-of-the art does not sufficiently address either issue: multiple concurrent, distributed personal agents would fall victim to the classical CAP theorem [74], and the current approach of using password based authentication as fallback to biometric (used on most current smart phones, tablets, laptops, and other mobile devices) in our scenario suffers from the lack of trusted user interfaces. We require new solutions to both.

# Section b: Methodology

Our main, novel approach is to create a direct mapping of user location and behavior between the physical and digital worlds (cf. Figure 1 in part B1). This will have a huge impact not only on trust and scalability issues (as our two main challenges in Digidow), but will act as an additional enabler for location-based services that currently rely on individual GPS tracking on users' mobile phones. That is, location-based services will become available to users without smart phones and without sacrificing privacy, which in itself is fundamentally novel.

In terms of scientific methodology, we will design methods and protocols, prototype these designs in the form of working implementations, and analyze the prototypes in simulation and laboratory experiments. Some of the tasks are inherently explorative research with the implication that analysis and review in the first iterations (cf. process described in section 17) will be limited until the problem space is understood sufficiently well.

## 12 Modeling individuals in associated personal agents: the Digital Shadow Model (DSM)

The fundamental enabling approach in this project is that personal agents (in the digital world) track their associated individuals (in the physical world) as closely as possible, and therefore create a direct mapping. We also say that their digital traces follow individuals, and therefore speak of a *digital shadow*. Upon closer consideration, this is however not a simple, unidirectional location mapping from the physical to the digital world. As individuals move through the physical world, their actions not only *create* digital traces, but are also *enabled* by digital services — including payment (e.g. for public transport), border control, physical access control, etc. The resultant mapping is therefore bidirectional and inherently probabilistic. The central, unified task in Digidow will be to create the *digital shadow model (DSM)* of the individual's time-dependent location and behavior hosted and estimated inside the personal agent. DSM will cover at least the following aspects:

- three-dimensional **location** of the individual in the physical world: We will start with a simple 3D model based on the global coordinate system as used by GPS and extend it with the accuracy of each location estimate. Depending on the sensor that contributed a location estimate, the assumed area of the individual's whereabouts can be differently sized (and shaped).
- (typical and atypical) **behavior** of the individual: Modeling user actions and distinguishing between typical/normal and atypical/abnormal states can be arbitrarily complex (even predicting typical contexts on mobile devices carried by users does not yield high accuracy [75]). We will therefore start with standard supervised classification of the individual's interactions (such as payment for various services) based on a time-limited training period assumed to be typical and therefore used as ground truth and then extend.

The combination of location and behavior describe the assumed state of an individual, which is internal to DSM. Both location and behavior are sensed in the form of external **events** generated in the physical world and used to update the internal model in the digital world. Every event carries at least two fields of meta data:

- **timestamp**: Events happen at a specific time, and only describe location or behavior at that moment. The longer an event is in the past, the more uncertain the current internal state of the model will be concerning that aspect. Decisions in the personal agent (e.g. if a payment transaction should be approved or which details of the individual's digital identity should be sent to an verifier) will be based on the recent history of events and an extrapolation (time series prediction) to estimate the state at the time of the interaction. [1]
- **probability**: No sensor is completely accurate, neither for measuring location nor for biometric identification. Especially location and identification/authentication events will therefore be annotated with an estimate of their accuracy based on the specific sensor that took the measurement. Ideally, trusted sensors will themselves provide this probability estimate, but the model should account for erroneous (or malicious) events and estimates. Detecting impossible (or very improbable) events by specific sensors can be used to lower their reputation and provide feedback to other personal agents concerning future estimates from that sensor.

The resulting DSM is multi-factorial (three dimensions for location, one for time, one or more for the behavior) and probabilistic. Updates of the internal assumed state of the individual based on external events will require an error model that takes into account typical behavior (work/home/shopping/commuting cycles) but supports atypical states (holidays, travel, etc.). In ubiquitous computing, a large body of prior work on modeling user behavior exists, including context awareness (with the PI's PhD thesis being just one example in this area), activity recognition, location tracking, and many others. This is one of the reasons why we see this proposal as being rooted in these scientific areas and why we expect publications to target that community in addition to specific contributions in the four areas detailed below.

However, there is a significant challenge in applying these existing results to modeling in the scope of Digidow: for personal agents to be reasonably secure and to enable their execution on cheap embedded hardware and/or at global scale of one agent per human, they need to be **as simple as possible**. Complexity leads to insecure code, and complex models may require significant memory and computing resources. Both effects are detrimental to the aims of Digidow, and therefore DSM itself should be simple (mathematically and in terms of memory and computational requirements) to potentially support formal validation for correctness of the model (and in extension of the code base of the personal agent). Additionally, it is more likely that simple models will be communicable in the sense that automatic decisions based on the model (e.g. if an authentication attempt or a payment transaction should be granted) can be explained to end-users (cf. section 3, part B1).

It is currently unclear if a single model will be able to capture all aspects and still be usable for decision making, or if multiple interacting models (such as one per type of biometric authentication sensor, one for each location event source, one for the estimated physical location, one for the behavior of the individual, etc.) will be a better approach. Designing DSM is one of the central tasks in project Digidow.

Independently of the model details, an important function is full logging of all external communication. One major advantage of personal agents is that they can keep a tamper-resistant log of all interactions of a single

[1] Note: Within the scope of Digidow, we assume that an absolute frame of reference for timestamps exists and explicitly discard relativistic, gravitational, or other effects that lead to time dilation in the physical world. We argue that this assumption is (currently) warranted for global (but Earth-bound) authentication use cases. Nonetheless, time synchronization between distributed hosts is a well-known non-trivial problem on practical networks. For Digidow, we will assume time synchronization protocols accurate enough for liveliness guarantees (NTP being the prime example) and will consider malicious tampering with timestamps and synchronization as part of the security analysis, but will not contribute to synchronization of distributed clocks itself.

individual visible in or assisted by the digital domain. Even if a verifier demands access to more data than seemingly necessary for a transaction (e.g. forcing customers to declare their names, addresses, and dates of birth to participate in a bonus program even if only a pseudonymous identifier derived from but unlinkable to the citizen identity number would be technically and legally required), the complete transfer and every future occurrence can be logged and retained for inspection and potential legal recourse. No verifier will be able to request data items in the background without this request being recorded in the log. We will investigate realistic options to require verifiers to specify reasons for requesting items of an individual's identity and if logging these claimed reasons will contribute to privacy compliance (Android versions $< 6.0$ are a counter-example where fine-grained information about requested permissions and user "consent" does not directly improve privacy).

Work in all four project areas will directly interact with this DSM, both contributing external influences (location tracking events, biometric identification data, and authentication requests from verifiers) and using the internal state as the basis for decisions (authentication based on the estimated location of the individual).

## 13 Area A: Authentication

**Face authentication** has the conceptual advantage of required infrastructure (cameras) often already being installed. The problem with existing infrastructure is that individual systems are each configured and calibrated differently, including illumination as well as different hardware sensors and lenses. Face authentication using different data sources is difficult, which is why cross-system face authentication is not yet considered fully solved. Our goal is to enable face authentication from different data sources. Deep learning provides the advantage of data preprocessing being embedded into more powerful models with the potential for increased authentication performance. We will investigate deep learning for face authentication in the context of building a cross-system approach with different sources of face data.

**Biometric template storage** in the personal agent is a basic requirement for matching live sensor data with the expected user template. However, this biometric template is the **most critical item** in the individual's data, as it cannot be changed after a compromise. Even applying secure coding practices, the possibility for data leaks remains (cf. the `openssl` Heartbleed vulnerability). We will therefore split storing and processing biometric templates from other parts in the personal agent, with the implication that the data processing pipeline will be split over multiple realms of trust. The goal is to minimize the possibility of leaking biometric data under the assumption that the personal agent is under various classes of attack concerning its internal memory model.

**Multi-party authentication protocols** are another new requirement caused by the Digidow architecture. As biometric authentication (measuring and initial data pre-processing in biometric sensors, feature extraction and matching in personal agents) leads to use of the digital identity at a verifier in the next step, some form of proof of authentication will need to be forwarded to a third party (the verifier). With privacy as one of our main goals, simply forwarding results of the biometric matching is impossible, and we will therefore investigate various cryptographic options for designing 3-party biometric authentication. We will start with zero knowledge proofs and blind signatures for proof-of-liveness on top of an Axolotl ratchet for privacy and forward secrecy guarantees as well-known methods, but are aware that even the problem space is not fully understood at the time of this writing. However, most current protocol designs are synchronous, and this has direct implications on the network service protocols to be developed in area D.

**Identity federation** is comparatively well-understood. The aim is to enable use of digital identity across a wide range of verifiers, and key points are therefore openness, interoperability, and scalability. U2F/UAF from the FIDO alliance, OpenID/OAuth as web login standards, and respective protocols from STORK and FutureID will be used as the basis for quantitative and qualitative analysis towards the end of project Digidow. We intend to implement multiple different identity federation protocols (those that can be used to embed our multi-party biometric authentication method) for maximum interoperability with existing services.

## 14 Area B: Tracking

As part of our research in area B, we will run user studies with voluntary participants to collect and analyze their location traces and behavior. For a more detailed description of the procedure to handle ethical aspects of this data collection, please refer to the respective annex documents (explanation of the procedure and draft form for informed consent declaration).

**A physical location model** will be the first building block for having an agent-centric representation of the individual. We will first design a simple but still flexible model that will translate external tracking events from manifold tracking resources (cf. section 8) to physical locations. We intend to use WGS84 georeferences as this might currently be seen as the lowest common denominator in terms of location representation. On top of the location model, we will develop services and interfaces available for trusted external tracking resources.

**Physical tracking and event handling** will combine the internal location representation with asynchronously appearing external tracking events. As a first step, we will make use of the WiFi based tracking infrastructure hosted at the JKU (cf. Smart Information Campus [76, 27]). In its current state, the tracking system still requires a WiFi transmitter (e.g. a smart phone) worn by the individual that broadcasts beacon signals on a continuous basis. Although this is not a DfP solution, the infrastructure is available campus-wide and yields adequate location updates valuable for evaluating tracking events in a real-life large scale setup.

In a second step we will deploy a VSN tracking system similar to the one presented in [77]. To that end, we aim for a cooperation with the authors that have already agreed to supporting the project. Although the VSN will provide tracking events on a smaller scale, we expect to gain valuable insights for concluding on large scale usage. In both approaches, the Digidow personal agent will support the tracking systems as it is aware of the potential follow-up whereabouts the individual might visit in the future. This will be represented by the behavior model. By supporting two sources of location events, we will develop insights on: coordinating concurrent tracking sources; seamless integration of tracking data asynchronously delivered by two or more tracking sources; integrating tracking results of arbitrary precision; scalability of VSN; and supporting tracking systems with a-priori knowledge on the person being tracked.

**The digital behavior model** is the counterpart to the location model that will hold probability distributions reflecting the potential future whereabouts and actions of the individual before the actual tracking or verifier interaction event has arrived. We will focus on heuristic models and machine learning techniques and investigate the behavior of individuals throughout their everyday life, especially when traveling, crossing borders, and using (public) transportation. From that we will derive our model that will also be in charge of judging the authenticity and trustworthiness of a tracking or authentication event from a source with lower reputation. Additionally, this phase will yield mechanisms to adapt the reputation of tracking systems or single sensors (used in the respective tasks in area C).

**Merged probabilistic models** is the targeted result from this last phase that also comprises the insights achieved in the previous tasks. For that purpose we will develop prototypes that allow us to compare the impact of the usage of several state-of-the-art probabilistic tracking algorithms, such as Hidden Markov Models [78], the Kalman family of filters (e.g. EKF, UKF) [79], Particle Filters [80] and Gaussian Process Latent Variable Models [81]. We aim at developing a solution that will fit our requirements in terms of a) adequate tracking accuracy, b) low latency, and c) robustness. As Digidow will actively update the physical location, the update frequency of external tracking systems is of minor concern to the project.

## 15 Area C: Trust

As mentioned previously, establishing sufficient trust in a decentralized system as proposed for Digidow is **one of the two fundamental challenges** we have to address throughout the project for both points of view (cf. section 3 in part B1: personal agents trusting biometric sensors, and verifiers trusting personal agents). Consequently, there are more different tasks in area C than in other areas.

A central aspect is enabling trust in the personal agents themselves, which will hold critical personal data and, if successfully attacked, would allow nearly complete identity theft. In area C, a major concern will therefore be to secure the code and execution environments of personal agents and biometric sensors — note that verifiers are not in the main focus, as they will realistically use existing infrastructure (such as payment schemes for public transport, citizen databases for border control, existing customer databases for digital or physical shops, etc.) and therefore need to be trusted for the various scenarios. We will take a multifaceted approach to improve security of networked code on different layers:

**Semi-formal specification** of functional and non-functional requirements is the first step of certified software development and will describe the main functionality of personal agents and their interaction with external

parties. This specification is the basis for all following steps, but will need to be continuously updated to reflect latest results from the other areas.

**Safe languages** in the sense of strong static type systems and validating compilers are the first building block for producing secure, reliable code. In the first phase, we will prototype core functionality of personal agents in multiple (pure and non-pure) functional languages including (but not limited to) Haskell, Rust, Scala, and Go. Languages will be selected and evaluated by the expressiveness of their type system (Haskell, Scala), compiler infrastructure (Haskell, Rust, Go), memory handling (with interesting new concepts such as in Rust), code size (Haskell), maintainability, and portability. Specific quantitative evaluation criteria will be (considering the same functionality implemented in/with each language/platform): lines of code, percentage of "unsafe" (imperative, interfacing with native system libraries) code parts, programmer time to first functional prototype (as an indication of language complexity and existing library maturity), portability to (number of) execution environments with a single code base, and (more qualitatively measured) estimated maintainability. An important side effect of these multiple implementations is that we will show, from an early stage, interoperability of our network protocol specifications for future implementations.

**Code compartmentalization** within personal agents (and, to a lesser extent, biometric sensors) is an important approach to mitigate potential security vulnerabilities that can arise even when safe programming languages are used for their implementation. We will analyze the constituent parts of an agent – called modules or components – and evaluate different approaches to intra-application sandboxing such as multiple processes/threads connected with IPC interfaces (as used e.g. by `openssh` and `postfix` as two of the more popular servers to use that technique), OS-supported memory compartmentalization (used e.g. by the `Chromium` webbrowser engine), or more general container approaches (e.g. Linux namespaces/containers such as used by `docker`). Currently known components will include (but not be limited to):

- in-memory handling of the DSM (cf. section 12)
- cryptographic key handling and derivation of pseudonyms, including secure private key storage
- network communication, including biometric authentication protocols and identity federation
- handling of events for location and behavior tracking and estimating updates to the in-memory model
- trusted boot, run-time integrity verification, and remote attestation

Our evaluation will include an analysis of potential vulnerabilities and exploit possibilities and how they can be mitigated by the various compartmentalization techniques. A particular challenge is cross-platform compartmentalization that does not depend on specific operating systems and can be used on a wide variety of hosting environments including embedded hardware CPUs.

**Unikernel VMs** are one approach to minimize the attack surface of virtual guests or embedded code. After selecting the safe language best suited for our purpose, we will develop methods to compile the personal agent as a unikernel to run on hypervisors with a minimum amount of required code, removing the requirement of a general-purpose guest operating system and most parts of the standard libraries. With this approach, we significantly limit the impact of standard vulnerabilities and therefore complement the use of a safe language and code compartmentalization. Current research on unikernels such as `ClickOS`, `Clive`, `HaLVM`, `MirageOS` are a basis for initial prototypes. Main challenges for the application of unikernels are (not only for Digidow, but more generally): combining unikernel compilation with code compartmentalization (compartments within a unikernel vs. one unikernel per compartment); compiling (and certifying) unikernel binaries for multiple different hypervisors to support cross-platform deployment; and change management of hypervisor interfaces.

**Formal code verification** is the most critical and challenging part to improve the current state-of-the-art in secure coding, and will consequently be tackled in later phases of project Digidow in the light of previous insights. Referring the formal specifications developed in the respective tasks described above, we will design new tools to (semi-) automatically verify correspondence of the specifications with the source code and potentially the compiled binary representation(s). We will initially start with similar methods used for the seL4 secure microkernel/hypervisor project [82] (translating a formal specification to Haskell, validating this representation, translating to C, and verifying the second representation). That is, we will focus on tools to automate the cross-validation of different representations, potentially with manually added annotations in some of those representations (most probably in the final source code to reference relevant parts of the specification).

**Dynamic execution verification** will be used in addition to (static) validation of source code. We will develop run-time verification methods to further compare system states with a corresponding run-time model (automatically or manually) derived from the initial specification. Note that static and dynamic verification are complementary: static verification is especially beneficial in terms of safe memory management and avoiding well-known classes of vulnerabilities, while dynamic verification is particularly important for state machines typically used in client and server parts implementing networks protocols. For networked code like personal agents, both aspects need to be combined.

The methods described above all intend to improve trust in the correctness of running code, and are therefore mostly important for enabling end users to trust that their own personal agent is running as intended. The second issue of trust in remote parties requires complementary methods to ensure security:

**Chain of trust** is a technique that ensures that only trusted software and hardware components are loaded and used from bottom up in the boot chain where each component verifies the validity of the next one. In our decentralized identity system we will extend methods from TPMs to establish a complete chain of trust on every used sensor and agent. The root-of-trust will be located in the lowest layers of the systems (hardware and firmware based) and verification will reach up to the software component that performs the final authentications (application layer). *Remote attestation* is the technique of the TCG specification to prove to an online service that the system runs in a trustworthy state (i.e. computed PCR values that are considered reliable). It uses special Attestation Identity Keys (AIK) that are signed by certificate authorities (CA). The problem with this technique is that potential traceability of the end-user when using these AIKs, which is addressed by *Direct Anonymous Attestation* (DAA). The downside of these two attestation techniques is that both only allow the verification of one counterpart at a time. In our scenario where sensors and agents need to mutually trust each other, these techniques would require additional unnecessary communication. We will design a novel (and potentially very unconventional) cryptographic protocol to **unify mutual remote attestation and biometric authentication crossing trust boundaries**. This unification of two protocols is necessary to provide sufficiently low latency for near-real-time authentication.

**Secure cloud hosts** are required to execute personal agents and potentially verifiers. In this smaller and strictly focused task, we will extend our developed methods to provide a chain of trust with an intermediary layer of hypervisors as used in practical cloud environments. The main challenge is integration of TPM/PCR certification up to the hypervisor layer with a hardware TPM and passing on the PCR chain into virtual TPMs provided for each guest instance. That is, multiple (virtual) TPMs (e.g. one per virtualized personal agent) need to be initialized with the single trust chain bootstrapped on the hardware and host hypervisor. We will develop protocols to support such an extension of the PCR chain and prototype on standard server hardware equipped with TPMs and running various standard hypervisors (the private cloud infrastructure for these experiments is already available at our Institute of Networks and Security).

**Reputation of sensors** allows integration of legacy sensing and tracking infrastructure into Digidow. Our preferred approach to integrating biometric sensors is to rely on secure models with a certified chain of trust and remote attestation (as described in task "chain of trust" above) to cryptographically verify their expected behavior, specifically not to leak raw biometric sensor data to third parties, to return live data, and not to selectively censor any events. However, large VSNs have already been deployed in many cities, and it is unlikely that these will be completely modernized in the near future. Under the assumption of a realistic attacker model, we will therefore develop a decentralized reputation system for personal agents to accumulate feedback on sensors coupled with short-lived certificates to provide "soft" trust for legacy VSNs.

## 16 Area D: Network (service discovery and communication)

For mutual discovery and communication of personal agents, biometric sensors, and verifiers for each transaction, we face a set of hard requirements: a) global scalability, b) decentralized architecture, c) privacy preserving, and d) low latency. As outlined above (in section 10), no currently deployed network protocols fulfill all these requirements at the same time. To cope with these requirements all together, we propose a protocol consisting of 3 consecutive phases of operation. Phase 1 is responsible for pre-selection in service discovery with a special focus on worldwide scalability. Phase 2 selects the specific counterpart for service use in a

privacy-preserving way (with less focus on scalability due to the previous pre-filtering). Phase 3 handles bidirectional service communication (i.e. all transactions within Digidow) using a multi-hop overlay network route and aiming for low-latency service usage.

Although the network protocol phases are strictly consecutive during run-time execution, tasks will partly overlap during the project to support iterations in requirements and protocol design between the phases. Concerning all phases, the fundamental question on which party should initiate the communication needs to be investigated. At the moment we believe the personal agent to be best suited for the initiator role because it holds the digital model of the individual and hence is able to contribute to geospatial filtering. As this choice has an impact on scalability, privacy and latency, we will elaborate on this question in detail.

**Scalable service discovery** Building upon our central concept of DSM (section 12), we will apply geospatial filtering as a fundamental concept to solve scalability and latency in service discovery. In a first phase, we will evaluate the suitability of DNS for geographically based services, in particular as a slowly changing index for discovering biometric sensor and verifier infrastructure.

**Private service discovery** is necessary to keep colluding verifiers and assumed global adversaries from uncovering real identities behind cryptographic pseudonyms. This includes hiding the IP addresses of personal agents, which would make de-anonymization by traffic correlation trivial. While phase 1 can be used to discover services in the geographic target region (and rely on k-anonymity if the region is sufficiently populated), discovering and using the final service needs to be done via overlay networks with a separate network addressing scheme. We will evaluate DHT as a well-understood approach for global, decentralized overlay routing.

Scalability of phases 1 and 2 will be evaluated together by quantitative analysis in a wide-scale simulation. Depending on the complexity (and therefore expected computational load) of network protocols for secure agents developed in areas A and C, we will either run the simulation on our local private cloud infrastructure at INS or rent instances on public clouds like Amazon or Google (cf. other costs in Table 2).

**Low-latency service use** We differentiate explicitly between service discovery and use. Cryptographic communication typically requires synchronous request/reply message types e.g. for live biometric authentication and remote attestation (cf. sections 13 and 15). A first prototype will use Tor [71] infrastructure as a specific, widely deployed onion routing network to provide low-latency bidirectional communication between personal agents and biometric sensors on one and verifiers on the other leg of the 3-way cryptographic protocols. A particular challenge is to bind the result of (asynchronous) private service discovery in phase 2 to (synchronous) service use in phase 3 without leaking network identities. To this end, we will evaluate using the `.onion` domain space for transaction-specific, one-time identifiers derived from pseudonyms using identity-based cryptography.

**Fail-over strategies** will be required for any realistic use of packet-switched global networks. Potential approaches include multiple instances and caching, but without invalidating liveness guarantees for biometric authentication. Specific requirements will heavily depend on the protocols developed in phases 1 to 3.

## 17 Scientific process

The scientific process for areas A–D will be iterative:

1. Systematic analysis of functional and non-functional requirements
2. Systematic (literature and existing prototypes) review of potential options to design a solution
3. Development of working code to study the relevant requirements
   a) for studying the viability of approaches to biometric authentication, tracking, and personal agent trust, we will use rapid prototyping to integrate existing approaches with new designs
   b) for studying scalability of network and cryptographic protocols, we will use simulations
4. Qualitative and quantitative evaluation (including user studies) concerning the overall goals, with recourse to previous steps if necessary: to step 1 when requirements turn out to be conflicting (cf. section 18 and Table 1), to step 2 when the current set of options is exhausted, or to step 3 when the current design seems qualitatively sound, but is lacking in a quantitative sense.

The duration of each iteration is not defined statically, but depends on current challenges in each of the areas. However, we will create a deliberate synchronization point for steps 3 and 4 once every project year to build a

| Area | Risk | Alternative |
|---|---|---|
| Area A | Biometric matching pipeline cannot be separated over trust boundaries | Raw biometric measurements need to be transmitted to the personal agent for matching, with the disadvantage of requiring stronger trust assumptions. |
| Area B | Merged probabilistic DSM becomes too complex for realistic implementation in personal agent | Multiple interacting partial models can be used to implement separate aspects of the DSM, with the probable implication that decision-making rules will be more complex by consulting multiple models. |
| Area C | No existing language/platform is sufficient in terms of the evaluation criteria | Stronger focus on mitigation techniques including more fine-grained compartmentalization and additional run-time verification — at the cost of more computational/memory overhead and lower performance during execution. |
| Area C | Combination of formal specification and realistic implementation of personal agent is too complex for current methods on formal validation | Partial validation of only core parts of the code in combination with extended static and dynamic code verification still go significantly beyond current state-of-the-art in secure code, but will not offer formal security proofs. This option therefore requires combination with organizational security measures such as external certification agencies in addition to the technical measures applied within the project (similar to current hardware/software certification standards like EAL). |
| Area D | Private, synchronous service use cannot be guaranteed | Cryptographic protocols for attestation and authentication need to be redesigned to support near-synchronous communication. |

Table 1: Main risks and potential alternatives

single prototype integrating all recent results from the four areas and to evaluate the overall architecture with regards to main project goals (cf. section 19).

## 18 Risks and alternatives

All areas and the associated tasks carry various risks. Especially the fundamental challenges of establishing trust in a decentralized system and global scalability require novel methods to be designed, prototyped and evaluated. We are confident that the overall concept of Digidow is the only workable solution given our set of goals, but are aware of significant risk in reaching some of the sub-goals described above. The currently perceived main risks and potential alternative courses of action are summarized in Table 1 as a starting point for deviations from the current, tentative project plan (as outlined in Figure 1).

## 19 Detailed work plan

Figure 1 gives a tentative overview of the length and order of the tasks described above. Every area has a dedicated area leader (post-doc researcher if possible), as described below. There will be no static assignment of PhD students to tasks, but the work plan balances load across all project years. Lighter colors indicate tasks for which the theoretical groundwork is well understood and which require research on evaluating different approaches and adapting them for the Digidow architecture. Darker colors indicate that more fundamental research is required before even all aspects of the problems can be fully understood, and which are therefore associated with higher risk to the overall project goals at the time of this writing.

At the end of each project year, we will build one consolidated prototype integrating all current results from the four areas. These will act as proof-of-concept demonstrators to allow easier communication of insights and directions to a wider audience, and will guide research questions in the subsequent year.

## 20 Collaborations and the intellectual environment at the host institution

Since insecure – and more generally incorrect – code is undoubtedly one of the biggest problems in current computer science, embedding the project at JKU guarantees there is a possibility for fruitful discussions and

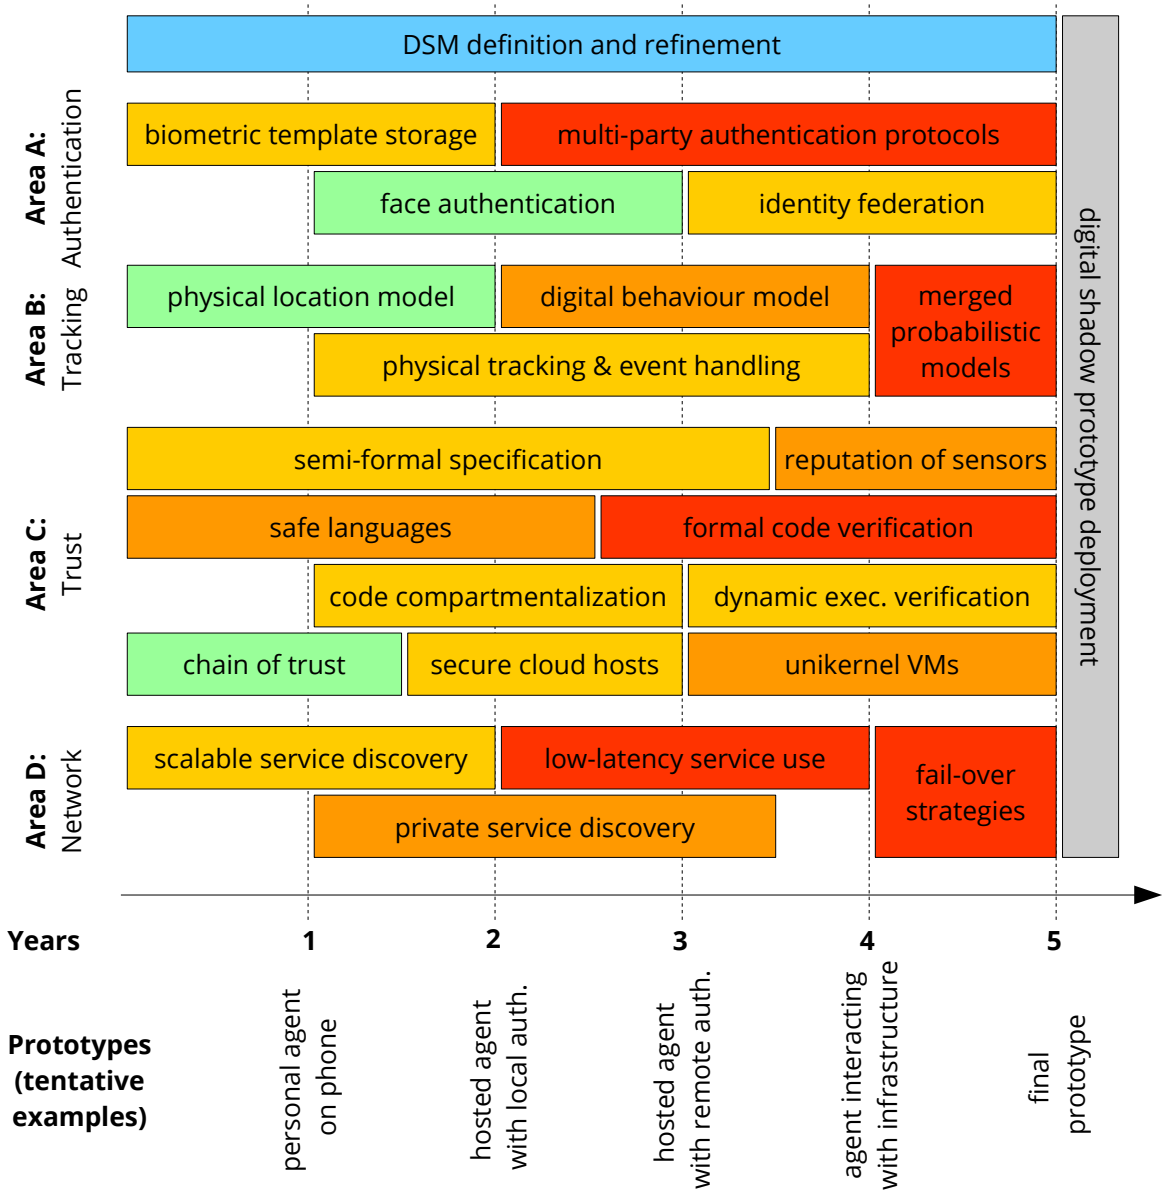


Figure 1: Tentative work plan, subject to change depending on results in respective previous tasks.

to acquire additional expertise if needed. "Secure Code" is a newly established strategic research focus at JKU, and the PI collaborates intensively with the groups of Prof. Armin Biere (Formal Verification, SAT solving technology), Prof. Alexander Egyed (Software Engineering with a focus on the impact of change), and Prof. Hanspeter Mössenböck (Systems Software with extensive practical experience with compilers and Java run-time environments). Additionally, there is a strong background on machine learning at JKU, including the Institute of Bioinformatics chaired by Prof. Sepp Hochreiter, a pioneer in deep learning techniques and the Department of Computational Perception chaired by Prof. Gerhard Widmer, currently on an ERC Advanced Grant on applying machine learning techniques to detecting emotion in music. Finally, the Institute of Pervasive Computing chaired by Prof. Alois Ferscha has extensive experience with many of the use cases in ubiquitous/pervasive computing that will be enabled or supported by Digidow. This combination of available outstanding expertise offers an excellent scientific environment for this project.

Specifically for visual sensor networks, the PI has already initiated a collaboration with Prof. Bernhard Rinner at the Alpen-Adria-University Klagenfurt, and will through this cooperation gain access to existing prototypes on camera systems equipped with TPMs, accelerating prototype work in areas B and C.

We explicitly note that these collaborators do not require financial resources from this project grant.

Concerning many aspects of creating and using digital ID, Austria has long experience with e-government in the form of the national citizen card (Bürgerkarte). The JRC u'smile lead by the PI in close coordination with INS is co-funded by the Austrian state printing house (Österreichische Staatsdruckerei), NXP Semiconductors, A1 Telekom Austria, and Drei-Banken-EDV under the lead of the PI. The PI is therefore already in close cooperation with the required technology (smart cards, person registers) and use case partners (mobile network operator, banks) for evaluating prototypes of digital identity systems, and leads the only consortium in Austria working on a digital (mobile phone) driving license. Together with the Finnish person register, we will create an intra-European standard for digital ID that will directly benefit from identity federation developed in Digidow.

## Section c: Resources (including project costs)

For this project, I request funding for a total of €1.992.750, mainly for personnel and costs associated with travel and publications. An overview of all costs and the total budget including overheads is shown in Table 2. The amounts written in brackets show the actual personnel costs that will be covered by the host institution and not requested to the project — respectively the total actual budget.

| **Cost Category** | | | ***Total in Euro*** | |
|---|---|---|---|---|
| **Direct Costs** | **Personell** | PI | (196 800) | |
| | | Senior Staff | (203 700) | |
| | | Postdocs | 681 600 | |
| | | PhD Students | 767 600 | |
| | | Other (Management/Technical) | (44 000) | |
| | *i. Total Direct Costs for Personnel (in Euro)* | | *(1 893 700)* | ***1 449 200*** |
| | **Travel Costs** | | 50 000 | |
| | **Equipment** | | 25 000 | |
| | **Other goods and services** | Consumables | 25 000 | |
| | | Publications | 37 500 | |
| | | Other (Audit) | 7 500 | |
| | *ii. Total Other Direct Costs (in Euro)* | | *145 000* | |
| **A - Total Direct Costs (i + ii) (in Euro)** | | | (2 038 700) | **1 594 200** |
| **B - Indirect Costs (overheads)** 25% of Direct Costs in Euro | | | (509 675) | 398 550 |
| C1 - Subcontracting Costs (no overheads) (in Euro) | | | | 0 |
| C2 - Other Direct Costs with no overheads (in Euro) | | | | 0 |
| **Total Estimated Eglible Costs (A + B + C)** (in Euro) | | | (2 548 375) | **1 992 750** |
| **Total Requested EU contribution** (in Euro) | | | | **1 992 750** |

| | |
|---|---|
| **Please indicate the duration of the project in months:** | **60** |
| **Please indicate the % of working time the PI dedicates to the project over the period of the grant:** | **40%** |
| **Please indicate the % of working time the PI spends in an EU Member State or Associated Country over the period of the grant:** | **90%** |

Table 2: Summary of all eligible costs and total project cost estimate, rounded to whole Hundreds.

**Personnel** As head of the Institute of Networks and Security (INS), research on digital identities is a strategic topic for the PI and the whole INS staff. 40% of my work time (which is nominally 40h/week) will be dedicated specifically to project Digidow for management and for providing personal research input. Due to the duties as head of the institute, I will be spending most of my time at JKU, but may take one sabbatical term (up to 6

months out of 60) for research leave at another academic institution, potentially outside the EU.

For Digidow, we will require 5 full-time researchers working exclusively on the project, with at least 2 out of 5 as post-doc to assist in managing areas A and D. Two staff members financed by JKU will contribute to managing areas B (Heinrich Schmitzberger, already post-doc) and C (Michael Hölzl, expected to finish his PhD shortly after the start of project Digidow). These are listed in the budget as contributing JKU staff without requesting funding for their employment. Other key members listed below and PhD students to be employed need to be funded by the Digidow budget.

The amount of funding per year is based on typical amounts for scientific staff in Austria (collective contract for Universities). Basic infrastructure and office space are provided by the host institution (JKU) and are paid through the overhead. Note that personnel costs have been based on the yearly wage index increase which is usually around 3% in Austria. We are also planning for one audit, which will be performed in order to guarantee formally correct cost statements. Additional "other" costs are for renting computational resources for large-scale simulations on service discovery protocols scalability.

**Travel and publications costs** The PI, post-docs, and PhD students are expected to – on average – travel to one international and one European conference per year to present results or participate in relevant workshop discussions. Travel costs are therefore estimated with €2.500 per person per year. Instead of conference travel in a specific year, some staff members may visit collaborators and their research groups for more direct exchange. Additionally, we expect costs for open access publications in the range of €2.500 per publication with an estimate of 3 publications per year.

**Equipment** The only requirement of hardware per person is a laptop computer and reasonable desktop equipment for people hired directly on the project. These are included mainly in the budget for the first years and are expected to be used throughout the project duration, with some upgrades in later years.

**Consumables** Most other equipment required for the scientific areas A–D (e.g. biometric sensors, embedded hardware boards for personal agent prototypes, setup equipment for user studies, current mobile devices as prototype hardware, etc.) is classified as consumables according to Austrian tax law.

### Project team

**René Mayrhofer** (head of the Institute of Networks and Security, INS) will act as principal investigator. He will contribute research input specifically to the DSM, embedded systems design, systems security, and cryptographic protocols, and will also be responsible for defining aims for the yearly project-wide prototypes and tracking associated cross-area integration efforts. In addition to project leadership outside his academic career, he currently has the same role in the Josef Ressel Center for User-friendly Secure Environments (JRC u'smile, funded until Sept. 2017) and is therefore experienced in managing a project group of comparable size.

**Heinrich Schmitzberger** is currently a post-doctoral researcher at INS and, having done his PhD on WiFi based location tracking [27, 28], will be working on the issues of areas B and D.

**Michael Hölzl** is a 3rd-year PhD candidate at the time of this writing, but is expected to hand in his PhD thesis shortly after the possible start of project Digidow. With his experience on cryptographic protocols, smartcard authentication [51, 54, 83, 84] and privacy-conscious digital identity, he will work mostly in area C and multi-party cryptographic authentication protocols in area A.

**Rainhard Findling** is a 3rd-year PhD candidate at the time of this writing, and is also expected to hand in his PhD thesis shortly after the possible start of project Digidow. Based on his experience with face authentication [85, 86] and other sensor based authentication methods [87, 88, 89], he will contribute to area A.

**Michael Sonntag** is the expert on Internet and data protection laws at INS. He will contribute to the analysis of legal aspects and their implications to Digidow, e.g. on cross-country identity federation protocols, certification of secure code and respective chains of trust, or requirements on storing biometric templates.

**Additional personnel** in the form of additional PhD students and/or post-doc researchers will be selected with the typical process of public, open job announcements on international web portals and mailing lists.

## References


[1] European Union, "Charter of the fundamental rights," *Official Journal of the European Communities*, no. 364/01, 2000. [Online]. Available: http://www.europarl.europa.eu/charter/pdf/text_en.pdf

[2] STORK Consortium, "Stork project." [Online]. Available: https://www.eid-stork.eu/

[3] FutureID Consortium, "FutureID project." [Online]. Available: http://futureid.eu/

[4] NewP@ss Consortium, "The NewP@ss project." [Online]. Available: http://newpass.av.it.pt/

[5] FastPass Consortium, "FastPass - a harmonized, modular reference system for all European automated border crossing points." [Online]. Available: https://www.fastpass-project.eu/

[6] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," Technical Report 07-49, University of Massachusetts, Amherst, Tech. Rep., 2007.

[7] G. Ramkumar and M. Manikandan, "Face Recognition-Survey," *International Journal of Advances in Science and Technology (IJAST)*, pp. 260–268, 2013.

[8] B. Weyrauch, B. Heisele, J. Huang, and V. Blanz, "Component-Based Face Recognition with 3D Morphable Models," in *Conference on Computer Vision and Pattern Recognition Workshop, 2004. (CVPRW '04)*, Jun. 2004, p. 85.

[9] Z. Riaz, A. Gilgiti, and S. Mirza, "Face recognition: a review and comparison of HMM, PCA, ICA and neural networks," in *E-Tech 2004*, Jul. 2004, pp. 41 – 46.

[10] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711–720, Jul. 1997.

[11] C. Liu and H. Wechsler, "Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition," *IEEE Transactions on Image Processing*, vol. 11, no. 4, pp. 467–476, apr 2002.

[12] X. Li, L. Wang, and E. Sung, "AdaBoost with SVM-based component classifiers," *Engineering Applications of Artificial Intelligence*, vol. 21, no. 5, pp. 785–795, Aug. 2008.

[13] P. J. Phillips, "Support Vector Machines Applied to Face Recognition," in *Neural Information Processing Systems*, M. I. Jordan, M. J. Kearns, and S. A. Solla, Eds., vol. 10, 1998, pp. 803–809.

[14] M. Zhou and H. Wei, "Face verification using Gabor wavelets and AdaBoost," in *Proc. 18th International Conference on Pattern Recognition, Vol 1*, Y. Y. Tang, S. P. Wang, G. Lorette, D. S. Yeung, and H. Yan, Eds. Los Alamitos: IEEE Computer Society, Aug. 2006, pp. 404–407.

[15] A. K. Jain, B. F. Klare, and A. Ross, "Guidelines for best practices in biometrics research," in *International Conference on Biometrics (ICB)*, vol. 8, Phuket, Thailand, May 2015.

[16] K. W. Bowyer, K. Chang, and P. Flynn, "A survey of approaches and challenges in 3D and multi-modal 3D + 2D face recognition," *Computer Vision and Image Understanding*, vol. 101, no. 1, pp. 1–15, 2006.

[17] A. F. Abate, M. Nappi, D. Riccio, and G. Sabatino, "2D and 3D face recognition: A survey," *Pattern Recognition Letters*, vol. 28, no. 14, pp. 1885–1906, Oct. 2007.

[18] Y. Sun, X. Wang, and X. Tang, "Deep Learning Face Representation from Predicting 10,000 Classes," in *Proc. 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2014, pp. 1891–1898.

[19] Y. Sun, Y. Chen, X. Wang, and X. Tang, "Deep Learning Face Representation by Joint Identification-Verification," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 1988–1996.

[20] D. Ciresan, U. Meier, and J. Schmidhuber, "Multi-column deep neural networks for image classification," in *Proc. 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2012, pp. 3642–3649.

[21] S. Isaacman, R. Becker, R. Cáceres, S. Kobourov, M. Martonosi, J. Rowland, and A. Varshavsky, "Identifying Important Places in People's Lives from Cellular Network Data," in *Pervasive Computing*, D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, K. Lyons, J. Hightower, and E. M. Huang, Eds. Springer Berlin Heidelberg, 2011, vol. 6696, pp. 133–151.

[22] M. Srivatsa and M. Hicks, "Deanonymizing mobility traces: using social network as a side-channel," in *Proc. ACM conference on Computer and communications security 2012 (CCS '12)*. ACM Press, 2012, p. 628.

[23] J. Krumm, "Inference Attacks on Location Tracks," in *Pervasive Computing*, ser. Lecture Notes in Computer Science, A. LaMarca, M. Langheinrich, and K. Truong, Eds. Springer Berlin Heidelberg, 1 2007, vol. 4480, pp. 127–143.

[24] Y. Morales and T. Tsubouchi, "DGPS, RTK-GPS and StarFire DGPS Performance Under Tree Shading Environments," in *Proc. 2007 IEEE International Conference on Integration Technology (ICIT '07)*, Mar. 2007, pp. 519–524.

[25] B. Sobhani, E. Paolini, A. Giorgetti, M. Mazzotti, and M. Chiani, "Target Tracking for UWB Multistatic Radar Sensor Networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 1, pp. 125–136, Feb. 2014.

[26] C. Rizos, G. Roberts, J. Barnes, and N. Gambale, "Experimental results of Locata: A high accuracy indoor positioning system," in *Proc. 2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE,

Sep. 2010, pp. 1–7.
[27] H. Schmitzberger and W. Narzt, "Campus-Wide Indoor Tracking Infrastructure," *International Journal On Advances in Networks and Services*, vol. 4, no. 1 and 2, pp. 138–148, 2011.
[28] ——, "Leveraging WLAN Infrastructure for Large-Scale Indoor Tracking," *Proc. 6th International Conference on Wireless and Mobile Communications 2010 (ICWMC)*, pp. 250–255, Sep. 2010.
[29] H. Zou, H. Wang, L. Xie, and Q.-S. Jia, "An RFID indoor positioning system by using weighted path loss and extreme learning machine," in *Proc. IEEE 1st International Conference on Cyber-Physical Systems, Networks, and Applications 2013 (CPSNA)*. IEEE, Aug. 2013, pp. 66–71.
[30] X.-Y. Lin, T.-W. Ho, C.-C. Fang, Z.-S. Yen, B.-J. Yang, and F. Lai, "A mobile indoor positioning system based on iBeacon technology," in *Proc. 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society 2015 (EMBC)*. IEEE, Aug. 2015, pp. 4970–4973.
[31] E. Deretey, M. T. Ahmed, J. A. Marshall, and M. Greenspan, "Visual indoor positioning with a single camera using PnP," in *Proc. 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society 2015 (EMBC)*. IEEE, Oct. 2015, pp. 1–9.
[32] L. Esterle, P. R. Lewis, X. Yao, and B. Rinner, "Socio-economic vision graph generation and handover in distributed smart camera networks," *ACM Transactions on Sensor Networks (TOSN)*, vol. 10, no. 2, pp. 1–24, 2014.
[33] R. Harle, "A Survey of Indoor Inertial Positioning Systems for Pedestrians," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1281–1293, 2013.
[34] K. Liu, X. Liu, L. Xie, and X. Li, "Towards accurate acoustic localization on a smartphone," in *Proc. IEEE INFOCOM 2013*. IEEE, Apr. 2013, pp. 495–499.
[35] S. Lynen, M. Achtelik, S. Weiss, M. Chli, and R. Siegwart, "A robust and modular multi-sensor fusion approach applied to MAV navigation," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems 2013 (IROS)*, Nov. 2013, pp. 3923–3929.
[36] M. B. Kjærgaard, "A Taxonomy for Radio Location Fingerprinting," in *Location- and Context-Awareness*, ser. Lecture Notes in Computer Science, J. Hightower, B. Schiele, and T. Strang, Eds. Springer Berlin Heidelberg, Jan. 2007, vol. 4718, pp. 139–156.
[37] M. Youssef, M. Mah, and A. Agrawala, "Challenges: device-free passive localization for wireless environments," in *Proc. 13th annual ACM international conference on Mobile computing and networking*. Montreal, Quebec, Canada: ACM, 2007, pp. 222–229.
[38] A. Saeed, A. Kosba, and M. Youssef, "Ichnaea: A Low-Overhead Robust WLAN Device-Free Passive Localization System," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 1, pp. 5–15, Feb. 2014.
[39] M. Seifeldin, A. Saeed, A. Kosba, A. El-Keyi, and M. Youssef, "Nuzzer: A Large-Scale Device-Free Passive Localization System for Wireless Environments," *IEEE Transactions on Mobile Computing*, vol. 12, no. 7, pp. 1321–1334, Jul. 2013.
[40] L. Schauer, M. Werner, and P. Marcus, "Estimating Crowd Densities and Pedestrian Flows Using Wi-Fi and Bluetooth," in *Proc. 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MOBIQUITOUS '14)*. ICST, 2014.
[41] S. Sigg, M. Scholz, Shuyu Shi, Yusheng Ji, and M. Beigl, "RF-Sensing of Activities from Non-Cooperative Subjects in Device-Free Recognition Systems Using Ambient and Local Signals," *IEEE Transactions on Mobile Computing*, vol. 13, no. 4, pp. 907–920, Apr. 2014.
[42] A. Wagner, J. Wright, A. Ganesh, Zihan Zhou, H. Mobahi, and Yi Ma, "Toward a Practical Face Recognition System: Robust Alignment and Illumination by Sparse Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 2, pp. 372–386, Feb. 2012.
[43] T. Winkler and B. Rinner, "Security and Privacy Protection in Visual Sensor Networks: A Survey," *ACM Computing Surveys (CSUR)*, vol. 47, no. 1, pp. 1–42, 2014.
[44] P. Kulkarni, D. Ganesan, P. Shenoy, and Q. Lu, "SensEye: a multi-tier camera sensor network," in *Proc. 13th annual ACM international conference on Multimedia*. Hilton, Singapore: ACM, 2005, pp. 229–238.
[45] A. Rowe, D. Goel, and R. Rajkumar, "FireFly Mosaic: A Vision-Enabled Wireless Sensor Networking System," in *Proc. 28th IEEE International Conference on Real-Time Systems Symposium 2007 (RTSS)*. IEEE, Dec. 2007, pp. 459–468.
[46] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua, "Multicamera People Tracking with a Probabilistic Occupancy Map," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 2, pp. 267–282, Feb. 2008.
[47] M. Song, D. Tao, and S. J. Maybank, "Sparse Camera Network for Visual Surveillance – A Comprehensive Survey," *CoRR*, vol. abs/1302.0446, 2013. [Online]. Available: http://arxiv.org/abs/1302.0446
[48] M. Langheinrich, R. Finn, V. Coroama, and D. Wright, "Quo Vadis Smart Surveillance? How Smart Technologies Combine and Challenge Democratic Oversight," in *Reloading Data Protection*. Springer, Jan. 2014, pp. 151–182.
[49] T. Winkler, A. Erdelyi, and B. Rinner, "TrustEYE.M4: Protecting the sensor – Not the camera," in *Proc. 11th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. IEEE, Aug. 2014, pp.

159–164.

[50] C. Cornelius, A. Kapadia, D. Kotz, D. Peebles, M. Shin, and N. Triandopoulos, "Anonysense: privacy-aware people-centric sensing," in *Proc. 6th international conference on Mobile systems, applications, and services*. Breckenridge, CO, USA: ACM, 2008, pp. 211–224.

[51] M. Hölzl, R. Mayrhofer, and M. Roland, "Requirements for an Open Ecosystem for Embedded Tamper Resistant Hardware on Mobile Devices," in *Proc. of International Conference on Advances in Mobile Computing & Multimedia (MoMM '13)*. New York, NY, USA: ACM, 2013, pp. 249:249–249:252.

[52] P. Bichsel, J. Camenisch, T. Groß, and V. Shoup, "Anonymous credentials on a standard java card," in *Proc. 16th ACM conference on Computer and communications security*, ser. CCS '09. New York, NY, USA: ACM, 2009, pp. 600–610.

[53] Trusted Computing Group, "Trusted Platform Module Library Specification, Family 2.0, Level 00, Revision .1.16," Oct. 2014. [Online]. Available: https://www.trustedcomputinggroup.org/resources/tpm_library_specification

[54] J. González, M. Hölzl, P. Riedl, P. Bonnet, and R. Mayrhofer, "A Practical Hardware-Assisted Approach to Customize Trusted Boot for Mobile Devices," in *Proc. Information Security Conference 2014 (ISC)*. Hong Kong: Springer International Publishing, 2014.

[55] T. Winkler and B. Rinner, "TrustCAM: Security and Privacy-Protection for an Embedded Smart Camera Based on Trusted Computing," in *Proc. Seventh IEEE International Conference on Advanced Video and Signal Based Surveillance 2010 (AVSS)*, Aug. 2010, pp. 593–600.

[56] S. Saroiu and A. Wolman, "I Am a Sensor, and I Approve This Message," in *Proc. Eleventh Workshop on Mobile Computing Systems & Applications*, ser. HotMobile '10. New York, NY, USA: ACM, 2010, pp. 37–42.

[57] P. Gilbert, L. P. Cox, J. Jung, and D. Wetherall, "Toward Trustworthy Mobile Sensing," in *Proc. Eleventh Workshop on Mobile Computing Systems & Applications*, ser. HotMobile '10. New York, NY, USA: ACM, 2010, pp. 31–36.

[58] T. Nyman, J.-E. Ekberg, and N. Asokan, "Citizen Electronic Identities using TPM 2.0," *arXiv:1409.1023 [cs]*, Sep. 2014, arXiv: 1409.1023. [Online]. Available: http://arxiv.org/abs/1409.1023

[59] FIDO Alliance, "FIDO UAF Protocol Specification v1.0," Dec. 2014. [Online]. Available: https://fidoalliance.org/specifications/download/

[60] ——, "Universal 2nd Factor (U2F) Overview v1.0," May 2015. [Online]. Available: https://fidoalliance.org/specifications/download/

[61] F. Zambonelli, N. R. Jennings, and M. Wooldridge, "Developing multiagent systems: The Gaia methodology," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 12, no. 3, pp. 317–370, 2003.

[62] A. Rowstron and P. Druschel, "Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems," in *Middleware 2001: IFIP/ACM International Conference on Distributed Systems Platforms Heidelberg 2001*, R. Guerraoui, Ed. Springer Berlin Heidelberg, 2001, pp. 329–350.

[63] B. Zhao, L. Huang, J. Stribling, S. Rhea, A. Joseph, and J. Kubiatowicz, "Tapestry: A Resilient Global-Scale Overlay for Service Deployment," *IEEE Journal on Selected Areas in Communications*, vol. 22, no. 1, pp. 41–53, Jan. 2004.

[64] K. Arabshian and H. Schulzrinne, "An ontology-based hierarchical peer-to-peer global service discovery system," *Journal of Ubiquitous Computing and Intelligence*, vol. 1, no. 2, pp. 133–144, 2007.

[65] Y. Doi, S. Wakayama, M. Ishiyama, S. Ozaki, and A. Inoue, "On Scalability of DHT-DNS Hybrid Naming System," in *Technologies for Advanced Heterogeneous Networks II: Proc. Second Asian Internet Engineering Conference, AINTEC 2006*, K. Cho and P. Jacquet, Eds. Springer Berlin Heidelberg, Nov. 2006, pp. 16–30.

[66] G. Pirrò, D. Talia, and P. Trunfio, "A DHT-based semantic overlay network for service discovery," *Future Generation Computer Systems*, vol. 28, no. 4, pp. 689–707, Apr. 2012.

[67] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2009. [Online]. Available: http://www.bitcoin.org/bitcoin.pdf

[68] C. Decker and R. Wattenhofer, "Information propagation in the Bitcoin network," *IEEE Thirteenth International Conference on Peer-to-Peer Computing (P2P), 2013*, pp. 1–10, Sep. 2013.

[69] I. Miers, C. Garman, M. Green, and A. D. Rubin, "Zerocoin: Anonymous Distributed E-Cash from Bitcoin," in *Proc. IEEE Symposium on Security and Privacy (SP), 2013*. IEEE, May 2013, pp. 397–411.

[70] L. Wang and J. Kangasharju, "Measuring large-scale distributed systems: case of BitTorrent Mainline DHT," in *IEEE Thirteenth International Conference on Peer-to-Peer Computing (P2P), 2013* , Sep. 2013, pp. 1–10.

[71] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proc. USENIX 2004*, 2004, pp. 303–320.

[72] A. Tran, N. Hopper, and Y. Kim, "Hashing It out in Public: Common Failure Modes of DHT-based Anonymity Schemes," in *Proc. 8th ACM Workshop on Privacy in the Electronic Society*, ser. WPES '09. ACM, 2009, pp. 71–80.

[73] Y. Bai, L. Di, and Y. Wei, "A taxonomy of geospatial services for global service discovery and interoperability," *Geoscience Knowledge Representation in Cyberinfrastructure*, vol. 35, no. 4, pp. 783–790, Apr. 2009.

[74] E. Brewer, "CAP Twelve Years Later: How the "Rules" Have Changed," *Computer*, vol. 45, no. 2, pp. 23–29, 2012.

[75] R. Mayrhofer, *An Architecture for Context Prediction*, ser. Schriften der Johannes-Kepler-Universität Linz. Trauner Verlag, Apr. 2005, vol. C 45.

[76] Institute of Business Informatics - Software Engineering, JKU Linz, "JKU Smart Information Campus." [Online]. Available: http://dg.jku.at/about.php?menuid=12

[77] B. Rinner, L. Esterle, J. Simonjan, G. Nebehay, R. Pflugfelder, G. Fernandez Dominguez, and P. R. Lewis, "Self-Aware and Self-Expressive Camera Networks," *Computer*, vol. 48, no. 7, pp. 21–28, Jul. 2015.

[78] J. Kosecka and F. Li, "Vision based topological Markov localization," in *Proc. IEEE International Conference on Robotics and Automation, 2004 (ICRA '04)*. IEEE, 2004, pp. 1481–1486 Vol.2.

[79] S. Y. Chen, "Kalman Filter for Robot Vision: A Survey," *IEEE Transactions on Industrial Electronics*, vol. 59, no. 11, pp. 4409–4420, Nov. 2012.

[80] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, Feb. 2002.

[81] B. Ferris, D. Fox, and N. D. Lawrence, "WiFi-SLAM Using Gaussian Process Latent Variable Models." in *IJCAI*, vol. 7, 2007, pp. 2480–2485.

[82] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood, "sel4: Formal verification of an os kernel," in *Proc. ACM 22nd Symposium on Operating Systems Principles (SOSP '09)*. New York, NY, USA: ACM, 2009, pp. 207–220.

[83] M. Hölzl, E. Asnake, R. Mayrhofer, and M. Roland, "Mobile Application to Java Card Applet Communication Using a Password-authenticated Secure Channel," in *Proc. 12th International Conference on Advances in Mobile Computing and Multimedia*, ser. MoMM '14. New York, NY, USA: ACM, 2014, pp. 147–156.

[84] ——, "A Password-authenticated Secure Channel for App to Java Card Applet Communication," *International Journal of Pervasive Computing and Communications (IJPCC)*, vol. 11, pp. 374–397, Oct. 2015.

[85] R. D. Findling, "Pan Shot Face Unlock: Towards Unlocking Personal Mobile Devices using Stereo Vision and Biometric Face Information from multiple Perspectives," Master's thesis, Department of Mobile Computing, School of Informatics, Communication and Media, University of Applied Sciences Upper Austria, Softwarepark 11, 4232 Hagenberg/Austria, Sep. 2013.

[86] R. D. Findling and R. Mayrhofer, "Towards Pan Shot Face Unlock: Using Biometric Face Information from Different Perspectives to Unlock Mobile Devices," *International Journal of Pervasive Computing and Communications*, vol. 9, no. 3, pp. 190–208, Sep. 2013.

[87] R. Mayrhofer, H. Hlavacs, and R. D. Findling, "Optimal Derotation of Shared Acceleration Time Series by Determining Relative Spatial Alignment," *International Journal of Pervasive Computing and Communications (IJPCC)*, vol. 11, no. 4, Oct. 2015.

[88] R. D. Findling, M. Muaaz, D. Hintze, and R. Mayrhofer, "ShakeUnlock: Securely Unlock Mobile Devices by Shaking them Together," in *Proc. MoMM 2014: 12th International Conference on Advances in Mobile Computing and Multimedia*. New York, NY, USA: ACM Press, December 2014, pp. 165–174.

[89] R. D. Findling and R. Mayrhofer, "Towards Device-to-User Authentication: Protecting Against Phishing Hardware by Ensuring Mobile Device Authenticity using Vibration Patterns," in *Proc. 14th International Conference on Mobile and Ubiquitous Multimedia (MUM'15)*. ACM, Dec. 2015, pp. 131–136.