A Privacy-aware Fall Detection System for Hospitals and Nursing Facilities

Erik Krempel* Pascal Birnstill** and Jürgen Beyerer***

Fraunhofer Institute of Optronics, System Technologies and Image Exploitation IOSB, Fraunhoferstr. 1, 76131 Karlsruhe (Germany)

* *erik.krempel@iosb.fraunhofer.de*

** *pascal.birnstill@iosb.fraunhofer.de*

*** *juergen.beyerer@iosb.fraunhofer.de*

09.12.2016

Abstract

Hospitals and nursing facilities are confronted with a critical shortage of qualified nursing staff. At the same time, patient safety must not fall by the wayside. Therefore we are facing a growing demand for assistance technologies that free up time for medical responsibilities. We introduce a fall detection system based on cameras and computer vision algorithms, which satisfies the high privacy demands of hospitals and nursing facilities. Our system explains its operations to patients and staff in order to establish transparency. Whenever we show video data to a nurse, it is either anonymized using image processing techniques or protected against misuse through usage control enforcement. We thus provide a system with a high level of functionality without sacrifying privacy.

I. Introduction

Societies in industrial countries are growing older and older. Since this leads to a shortage of qualified nursing staff, more and more technologies are developed in order to preserve patient safety and particularly to protect the elderly. Fall detection is a patient safety task, which is addressed by such new technologies, and which is particularly relevant for hospitals and nursing homes. In addition to the safety-related functionality, technologies to be deployed in such sensitive environments demand high standards concerning privacy and information security. Next to respecting the privacy of those affected, achieving acceptance among all stakeholders necessarily requires a system design that makes its internal procedures transparent.

In this work, we introduce a prototype of a video-based fall detection system with a strong emphasis on privacy, data protection requirements, and transparency. Starting from the requirements of all relevant stakeholder groups, we outline our system design and particularly elaborate on how privacy-related requirements are realized and enforced.

Relying on cameras in this context may seem contradictory, however, according to our requirements analysis it is the most suitable approach, since it allows a first-level situation assessment of potential emergencies from a distance. Furthermore, such systems can be deployed without major changes to the building infrastructure and, compared to sensors worn on the body, their protective effect is not limited to specific groups.

The task of fall detection is realized by means of a computer vision algorithm. Upon detecting a potential fall, it alerts the nursing staff via mobile devices, e.g., smart phones. The alerting procedure unobtrusively integrates into the daily routines of nurses so that a dedicated human operator is not required.

In terms of privacy protection we apply privacy filters (anonymization techniques) on video data to be released for first-level assessment. As privacy filtering of video data for protecting the identities of captured persons is the subject of works published by the authors, we only briefly outline how the application of such mechanisms is enforced in our system design. We also prevent a redistribution of video data from the nursing staff's mobile devices by employing distributed usage control technology augmented with information flow tracking as introduced by Pretschner et al.[1] as well as Hilty et al.[2]

Our system works transparently: It explains its state of operation to patients, visitors, and staff members using dedicated screens on the monitored corridors and by this means contributes to confidence and acceptance among all stakeholders.

This work is structured as follows. After discussing related work in terms of existing concepts for detecting falls in Section II, we explain the requirements of the relevant stakeholder groups in Section III. Section IV introduces our system design, while Section V elaborates on the mechanisms required to enforce the privacy-related requirements. Finally, we conclude in Section VI.

---

1 A. Pretschner, M. Hilty, and D. A. Basin. Distributed Usage Control. Commun. ACM, 49(9):39–44, 2006

2 M. Hilty, A. Pretschner, D. A. Basin, C. Schaefer, and T. Walter. A Policy Language for Distributed Usage Control. In J. Biskup and J. Lopez, editors, ESORICS, volume 4734 of Lecture Notes in Computer Science, pages 531–546. Springer, 2007.

II. Related Work

Fall detection is an area of intense research driven by three distinctive concepts [3].

The first concept is based on sensors that the monitored persons carry on their body. Fast accelerations or certain movement patterns raise an alarm. As only individuals that are equipped with a sensor can be monitored, this solution is not optimal for scenarios such as hospitals or nursing facilities. Furthermore, additional technology is required for locating a potential emergency.

The second concept uses sensors that are either embedded in the floor or into special carpets in the monitored area. Once people fall the sensors detect the difference in weight or electromagnetic patterns and raise an alarm. The deployment of these sensors into already existing infrastructures is costly and therefore no ideal solution for our scenario.

Most importantly, both technologies lack the ability to remotely assess the situation. When an alarm is raised the assessment has to be done by humans in the area, which does not solve the problem given the aforementioned shortage of personnel.

The third technology applicable for detecting falls is computer vision. Video analysis algorithms process the images of video sensors searching for characteristic patterns of persons falling. In order to improve the detection rate, many fall detection algorithms work on 3d data, i.e., requiring range imaging cameras, or combine video data with other sensors. Video-based systems are the most suitable approach for the considered scenario even though they also exhibit specific drawbacks. Algorithms for video-based fall detection cannot be operated without a certain rate of classification errors (cf. Section IV) and require a considerable amount of computation power. In addition, video cameras raise many data protection issues that have to be solved, which constitutes the main objective of this work.

III. Requirements

The system design introduced in this work is strongly driven by stakeholder needs. In our hospital scenario four different stakeholder groups have to be considered: hospital operators, hospital staff, patients and legal stakeholders. We interviewed representatives of all four groups in order to collect their specific requirements concerning the envisioned system.

According to hospital operators the safety monitoring in patients' rooms is guaranteed by already existing processes such as regular patrol rounds of the nursing staff and emergency buttons on every bed and in the bathrooms. In contrast the situation is problematic on corridors between the wards and on

[3] N. Noury, A. Fleury, P. Rumeau, A. Bourke, G. Laighin, V. Rialle, and J. Lundy. Fall Detection - Principles and Methods. In Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE, pages 1663–1666, Aug 2007

the paths the patients take when they leave their rooms to get outside. Especially in the late evening and at night emergencies may remain undetected for critical amounts of time. As qualified nursing staff is already a scarce resource and usually occupied with medical tasks in the wards, it is not possible to increase patrol rounds in those areas. Moreover, as security personnel does not have the required medical expertise to decide how to react to emergencies only nursing staff can serve as first aiders. Therefore the system shall help to detect falls in remote areas [**req01**] and to appoint the nursing staff more efficiently [**req02**].

The requirements of the nursing staff are largely similar to those of the operators. The system should be easy to use and integrate unobtrusively in their daily tasks [**req03**]. Interaction with the system can only be tolerated in rare occasions. As medical staff may be occupied with emergency procedures, the system must not interrupt them during ongoing tasks [**req04**]. Furthermore, the system must allow the nursing staff to remotely assess the situation, i.e., to either decide on the needed medical equipment [**req05**] or to dismiss false alarms. Privacy protection is also an urgent demand of the nursing staff. As they are almost permanently facing the cameras in their working environment, they demand measures protecting them against surveillance in the workplace [**req06**].

Another set of requirements is of course introduced by the patients. In the first place, they also emphasize a high demand for privacy protection. Not only is it important to protect the patients against permanent surveillance [**req07**], but also to understand the impact of collected video data on the patients' dignity. Leaking a video of a patient in an emergency situation would have a devastating effect on the acceptance of the system. Therefore mechanisms have to be in place that prevent the leakage of video data [**req08**]. Being in an unfamiliar surrounding usually means a lot of stress for patients. A permanent surveillance by video cameras may even exaggerate this issue. Therefore a high level of transparency concerning the operation of the system and its data processing steps is required in order to compensate for the presence of the cameras [**req09**].

The final set of requirements is derived from applicable law, i.e., from legal stakeholders. The German data protection law has been used as a basis for the analysis of legal requirements. As Germany has very strict data protection regulations, a system in compliance with these regulations will most likely also be compliant with the national regulations of many other states. Legal requirements for fall detection systems based on video analysis have been analyzed in earlier collaborations of the authors for both, the German Federal Data Protection Act (BDSG) as well as the draft for the General Data Protection Regulation (GDPR) of the European Union[4].

In Germany, video surveillance monitoring the workplaces of employees is only allowed in exceptional cases. As this is by no means the purpose of the system to be designed, countermeasures against this kind of misuse have to

[4] Birnstill, P.; Bretthauer, S.; Greiner, S. & Krempel, E.
Privacy-preserving Surveillance: an Interdisciplinary Approach
International Data Privacy Law, Oxford University Press, **2015**

be taken. This legal requirement is equivalent to [**req06**] and is therefore not listed as a new requirement. Section 6 a of the Federal Data Protection Act (BDSG) is also applicable to such systems. It demands that information processing systems must not take automated individual decisions entailing legal or other adverse consequences for the person(s) affected. Thus, the system has to be designed in such a way that each decision is supervised by a human before it affects an observed person [**req10**]. Section 3 a of the Federal Data Protection Act (BDSG) demands that the amount of data being collected and processed by a video surveillance system must be minimized. Therefore the system must fulfill the principle of data economy [**req11**]. According to section 6 b, subsection 5 of the Federal Data Protection Act (BDSG) all collected video data has to be deleted as soon as possible after having served its purpose [**req12**].

This collection of stakeholder requirements leads to a list of 12 aspects that have to be addressed by the system design:

req01: Detect falls in remote areas

req02: Appoint nursing staff more efficiently

req03: Integrate into nursing staff's daily routines

req04: Do not interrupt nursing staff during ongoing tasks

req05: Allow for remote assessment of potential emergencies

req06: Prevent monitoring of employees

req07: Prevent permanent surveillance of patients

req08: Prevent leakage of video data

req09: Provide a high level of transparency

req10: Avoid automated individual decisions taken by the system

req11: Adhere to the principle of data economy

req12: Delete any collected data as soon as possible

IV. System Design

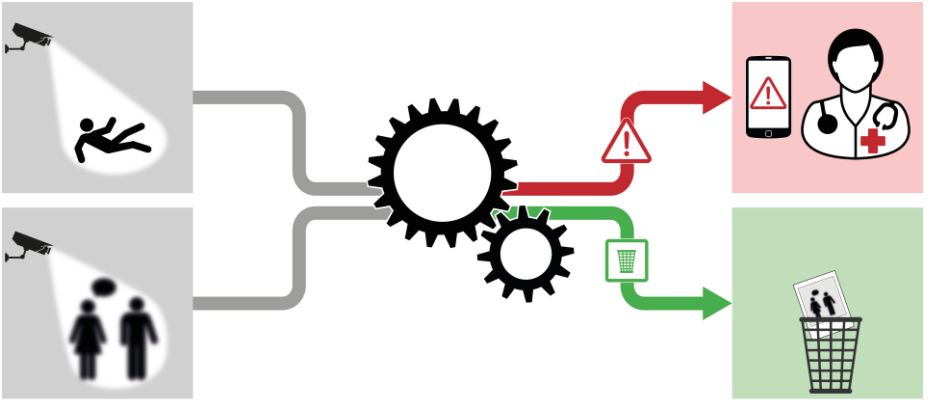

Figure 1: Workflow and data processing of the system

In order to incorporate all stakeholder requirements, we developed a new system design. Inspired by works of Roßnagel et al.[5] as well as Hornung and Desoi[6], we envision a privacy-aware smart video surveillance system design that separates functionality into distinct operational modes with different levels of privacy protection. A system's *default mode* is optimized for privacy: It collects and reveals a minimal amount of data. Event-specific *assessment modes* create views of the scene and available metadata, such that human operators can distinguish critical incidents from false alarms. At this stage, observed people's privacy is still protected as far as possible, typically by applying anonymization techniques before exposing video data. Finally, *investigation modes* unlock additional functionality for handling a specific type of incident and may involve deeper privacy intrusions. The rationale behind such system designs is to prevent groundless and unjustified intrusions into people's privacy. The operational modes are applied on the level of the particular cameras of the system: Typically, most of the cameras will be operated in the default mode, while only few cameras are involved into assessments or investigations of alarms.

We propose a system design that does not only influence how operators interact with the system, but which also makes its data processing steps transparent for the patients: For each camera at least one display is deployed in the proximity in order to inform the patients about the current operational mode as well as the corresponding implications on privacy.

---

[5] A. Roßnagel, M. Desoi, and G. Hornung. Gestufte Kontrolle bei Videoüberwachungsanlagen - ein Drei-Stufen-Modell als Vorschlag zur Grundrechtsschonenden Gestaltung. Datenschutz und Datensicherheit, 35(10):694–701, 2011.

[6] G. Hornung and M. Desoi. "Smart Cameras" und automatische Verhaltensanalyse. Verfassungs-und datenschutzrechtliche Probleme der nächsten Generation der Videoüberwachung. Kommunikation & Recht, 3:153–158, 2011.

IV.1. Default mode: Fall Detection

The system enters the default mode upon activation. Most of the cameras will be operated in the default mode for most of the time. As illustrated in Figure 1 video data is processed by a fall detection algorithm. As long as no potential fall has been detected, all data is deleted right after processing. To be more specific, an image is not stored for longer than 100 milliseconds in the memory of the application [partly satisfies **req11**]. In this mode, no human operator or user can get access to the video data. The screens in the observed area notify patients and staff that the system is activated, that data is currently being processed by the fall detection algorithm, and that nobody is granted access to the video data [partly satisfies **req09** + **req11**]. If and only if an algorithm detects a critical event, i.e., a person falling to the floor, the system changes into the assessment mode for the according camera [satisfies **req06** + **req07**].

IV.2. Assessment Mode: Anonymization of Video Data

Algorithmic fall detection is not yet an entirely solved problem in the area of video analysis. Classification errors of such algorithms cannot be eliminated completely. We distinguish between false positive errors, where a certain event, i.e., a fall, has been detected, but did not actually happen and false negative errors, where an event actually happened, but has not been detected. Since missing a fall may have severe consequences and hence would be devastating for the patients' and staff's confidence in the system, algorithms have to be parameterized conservatively: We minimize false negatives on the cost of a certain rate of false positive detections. As a consequence, the system design needs to cope with a certain amount of false alarms without privacy impact.

To achieve this, we design the system's operational workflow as follows. Upon detecting a critical event in the video stream of a camera, the system enters the assessment mode for this specific camera. The associated screen on the monitored corridor accordingly indicates the change of the system's operational mode accordingly: It announces that an emergency has been detected and that a nurse will evaluate this event shortly [partly satisfies **req09**].

At the same time, the system broadcasts an alarm to the mobile devices of the nursing staff in the proximity [partly satisfies **req02**]. In case none of the alerted nurses is able to accept the alarm, the systems sends further alarms while expanding the perimeter of recipients [satisfies **req04**]. Upon accepting an alarm, the respective nurse is associated to the assessment mode of the according camera. On all other mobile devices, the alarm is canceled.

The system provides the responsible nurse with an anonymized live view of the scene streamed to the mobile device. By means of user authentication on the mobile devices, the system can distinguish the nursing staff members and thus ensure that only a single nurse can get access to a camera being operated in assessment mode [partly satisfies **req08**, completes **req11**]. This is the first time that the system releases video data and access is restricted to a single member of the nursing staff for each alarm. All communication between the back-end and the mobile devices is sent via encrypted channels

in order to prevent attackers from accessing or modifying transferred data [partly satisfies **req08**]. Furthermore, we protect transmitted video data with distributed usage control technology (cf. Sec. V.1) in order to prevent misuse after access has been granted to a nurse [completes **req08**]. The quality of anonymized video data is sufficient to allow the nurse to distinguish actual emergencies from false alarms (cf. Sec. V.2). Only if the alarm is confirmed in this step, the systems changes into the investigation mode for the given camera.

This workflow is in line with the ban on automated individual decisions of information processing systems according to German data protection law. The system itself does not decide whether a person has fallen, but alerts members of the nursing staff to take over the responsibility to assess each potential emergency [satisfies **req10**].

IV.3. Investigation Mode: Bidirectional Communications Channel

When the system changes into the investigation mode for a detected event, the responsible nurse has already confirmed that the event is an actual emergency. In this mode, the system provides additional functionality in order to enable fast and effective aid [completes **req02**]. The responsible nurse is granted access to the unmodified video stream of the associated camera. This allows for more accurate situation assessment, which is necessary in order to decide which medical equipment is required for handling the emergency [satisfies **req05**].

As soon as the nurse accesses the video stream, the display on the accorrding camera indicates the switch into investigation mode. Moreover, it displays a live video channel to the responsible nurse who is currently viewing the video stream of the camera showing the patient. For this purpose, our system activates the front camera of the nurse's mobile device. Thus, in addition to increasing the transparency of the access to (non-anonymized) video data [completes **req09**], we enable bidirectional communication between the nurse and the patient. By this means, the nurse can reassure the patient that help is on its way and try to calm the patient.

Once the medical emergency has been handled, the responsible nurse closes the alarm, which puts the system back into its default mode. All collected data is deleted [satisfies **req12**] and access to the video stream is withdrawn.

V. Enforcing Privacy-related Requirements

Illnesses and hospital stays are very personal circumstances. Video data released by our system may even show persons in emergency situations. Hence, we are obviously handling particularly sensitive personal data, which must be protected by appropriate mechanisms. In the assessment mode (cf. Sec. IV.2), we release video data in an anonymized fashion. Accordingly, we have to ensure that an anonymization technique is applied before video data is sent to the mobile device of a nursing staff member. In the investigation mode, we even provide the responsible nurse with unmodified video data, for which it is particularly important to ensure that it is neither captured nor

redistributed from the mobile device. Our system enforces such requirements using *distributed usage control* technology.

V.1. Distributed Usage Control

Distributed Usage control (DUC) generalizes access control to the time after the initial access to data has been granted: It addresses obligations regarding the future usage of data, particularly in distributed settings[7]. Usage control requirements include rights and duties, e.g., “data may not be forwarded”, “data usage must be logged”, “data must be deleted after thirty days”. DUC requirements are specified in policies based on trigger events. We employ the *Obligation Sepcification Language (OSL)* introduced by Hilty et al.[8] for this purpose. Events are intercepted or observed by so-called *policy enforcement points (PEP)* as illustrated in Figure 2. PEPs forward events to a *policy decision point (PDP)*, which evaluates them against policies and replies with an *authorization action*, such as *allow, modify, inhibit*, and *delay*. In addition the PDP may trigger so-called *execute actions*, for instance sending a notification, writing a message to a log file, invoicing billing, and also deploying policies on other machines.

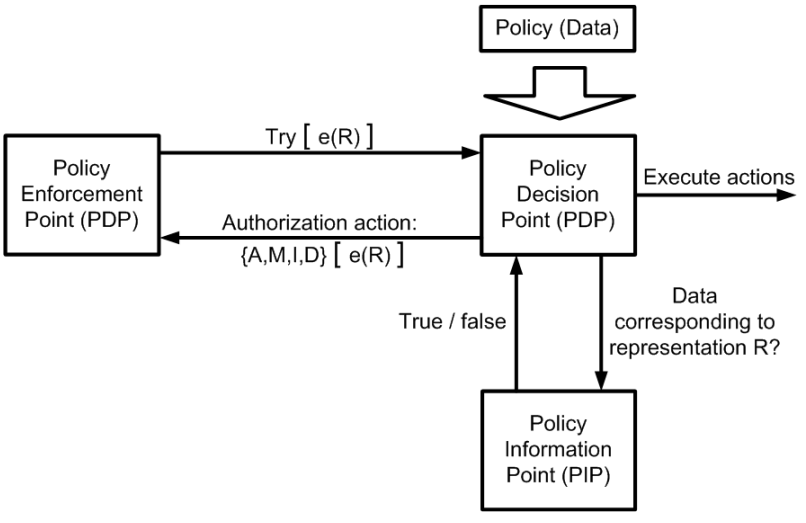

Figure 2: Generic usage control architecture with information flow tracking

In our system, “in the assessment mode, video data must be anonymized before release” is a typical usage control requirement that can be specified using the authorization action *modify*, which requires the execute action *obfuscate image* to be applied as long as video data is being streamed to the mobile device.

---

[7] A. Pretschner, M. Hilty, and D. A. Basin. Distributed Usage Control. Commun. ACM, 49(9):39–44, 2006.

[8] M. Hilty, A. Pretschner, D. A. Basin, C. Schaefer, and T. Walter. A Policy Language for Distributed Usage Control. In J. Biskup and J. Lopez, editors, ESORICS, volume 4734 of Lecture Notes in Computer Science, pages 531–546. Springer, 2007

In distributed settings, e.g., forwarding a data item with an attached policy to another system, DUC requirements can be enforced on the receiver's machine, too, requiring usage control enforcement mechanisms at the receiving end[9].

Because data usually comes in different representations – an image can be a pixmap, a file, reside in process memory, in a window, or be aggregated into an abstracted set of objects shown on the image – UC mechanisms have been augmented by information flow tracking technology [10,11]. One can then specify policies not only for specific fixed representations of data, but also on *all* representations of that data. These representations are tracked by information flow detection components. Policies then do not need to rely on events but can forbid specific representations to be created, also in a distributed setting. In other words, information flow tracking aims to answer the question into which representation within the (distributed) system the monitored data has been propagated.

Information flows are also detected by means of intercepting and interpreting events in the control flow of the system. The so-called *policy information point (PIP)* holds and interprets the information flow semantics of events and accordingly keeps track of new representations of data being created and of information flows between containers, such as files, windows, processes, etc. By this means, when evaluating an event concerning a data unit, the PDP can ask the PIP whether this data unit is a representation of a protected data unit, for which a policy must be enforced (cf. Figure 2). We use information flow tracking technology in order to protect video data released for fall assessments and investigations on the nursing staff's mobile devices against redistribution.

Reliability of distributed usage control enforcement and likewise the obtained level of security is based on the following assumptions. The integrity of policies and components of the UC infrastructure is ensured. The UC infrastructure is up and running and not tampered with, i.e., users do not have administrative privileges on their devices. Users also do not have the ability to suppress the communication between components of the UC infrastructure.

V.2. Enforcing Anonymization

We decouple image exploitation functionality into individual plugins that operate as a sequential plugin chain including a fall detection algorithm as well as an anonymization plugin, which implements the requirement [**req6 +**

[9] F. Kelbert and A. Pretschner. Data Usage Control Enforcement in Distributed Systems. In Proc. CODASPY, pages 71–82, 2013.

[10] M. Harvan and A. Pretschner. State-based Usage Control Enforcement with Data Flow Tracking Using System Call Interposition. In 3rd Intl. Conf. on Network and System Security, pages 373 –380, 2009.

[11] A. Pretschner, E. Lovat, and M. Büchler. Representation-independent Data Usage Control. In Data Privacy Management and Autonomous Spontaneus Security - 6th International Workshop, DPM 2011, and 4th International Workshop, SETOP 2011, Leuven, Belgium, September 15-16, 2011, Revised Selected Papers, pages 122–140, 2011.

req07], i.e., protecting patients and staff against permanent surveillance. The camera's video stream is the input of the plugin chain. The output is defined by the final plugin; in our case this is the anonymization plugin. The anonymization plugin can be configured to apply various image filters and pixel operations, such as Gaussian blurring and pixelization, on given regions of interest. By this means, we obfuscate observed people and thus protect their identities. This anonymization plugin also acts as a PEP, which is connected to the usage control infrastructure and provides the execute actions *delete image*, *obfuscate image*, as well as *pass through image*.

As long as our system operates in default mode, a dedicated policy demands that the anonymization plugin executes *delete image*, i.e., video data is neither stored nor released. If the system enters the assessment mode upon a potential fall has been detected, this state change is observed by the usage control enforcement infrastructure: Another policy applies for this operational mode and enforces the plugin to execute *anonymize image*. Finally, in case of a confirmed emergency, a third policy matches on the state change into the investigation mode and demands that the plugin executes *pass through image* in order to release unmodified images.

V.3. Protecting Video Data on Mobile Devices

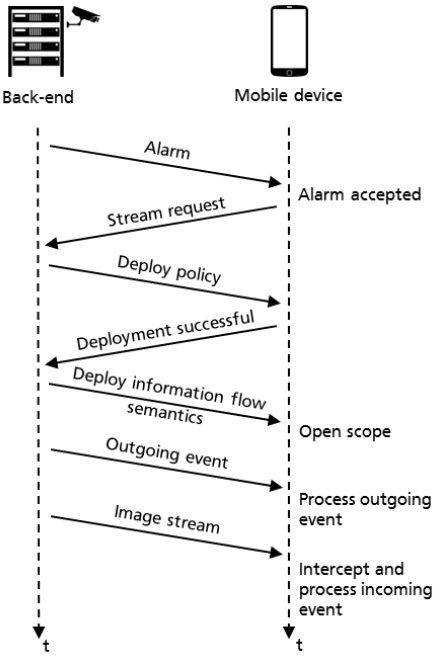

Figure 3: Protocol for cross-system information flow tracking

In order to prevent illegitimate redistribution of video data [**req08**], which is streamed to the mobile device of the responsible nurse, we have to (i) deploy an according policy on the usage control infrastructure of the mobile device, (ii) detect data of the video stream when received on the mobile device, and (iii) monitor this new representation of the video data in order to inhibit further representations of the data to be created. We achieve (ii) and (iii) by means of usage control combined with information flow tracking technology as outlined in Sec. V.1.

Figure 3 illustrates the particular protocol steps of processing such cross-system information flows. Streaming video data to a mobile device is triggered once a nurse accepts an alarm concerning a potential emergency. Accepting an alarm means that the authenticated nurse is associated to this alarm, i.e., no other staff member is able to connect to the corresponding video stream. Accepting an alarm further means that the nurse's mobile device requests the video stream from the back end of the system. The back-end generates a unique *dataID* for the video stream and triggers the deployment of a usage control policy on the mobile device's PDP, which refers to the dataID and governs restrictions concerning the video data. In this case, the policy demands that no further representations of the protected video data must be created, which includes that the video data must not be saved on the mobile device and also that the screen must not be captured while an application in the foreground has access to the video data. The unique dataID represents the data of the video stream within our information flow tracking infrastructure, i.e., within PIPs of the back-end and the mobile device.

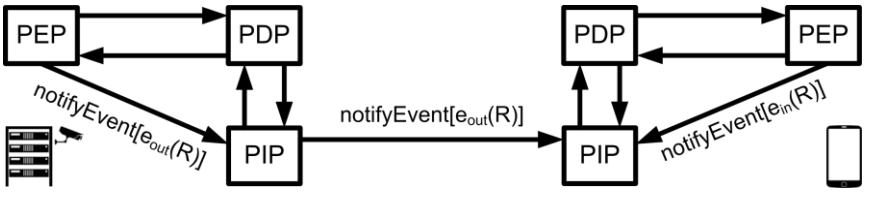

Figure 4: Cross-system information flow signaling in the DUC infrastructure

Enforcing this policy on the mobile device requires that its usage control infrastructure recognizes the incoming data as a new representation of the protected video stream. In order to establish this cross-system state for signaling the imminent information flow, an event e_{out} intercepted by the back-end's PEP for observing outgoing flows has to be interpreted by the mobile device's PIP (cf. Figure 4). For this, we deploy an information flow semantics specification on the mobile device's PIP, if not yet done. This semantics is required to establish the connecting link between the video data to be transferred and the already deployed policy: When processing the outgoing event forwarded from the back-end given the semantics, the mobile device's PIP opens a so-called *scope*, i.e., it prepares for incoming data over a network connection specified in the event (e.g., identified via IP address of the back-end as well as TCP source and destination port).

From the perspective of the mobile device, the incoming information flow is again indicated by means of events (cf. Figure 4, e_{in}). These events are detected by the mobile device's PEP, forwarded to its PIP, and tell the PIP

that data from a specific network connection is received by a certain process, i.e., in our case either the mobile app of our fall detection system, or any other media player application on the mobile device. Due to the scope that we established for the video stream, the PIP is able to associate the received data to the dataID given in the policy that we deployed on the mobile device's PDP earlier: The application accessing the video data constitutes the only permitted representation of the video stream on the mobile device.

In terms of enforcing our policy to inhibit the redistribution of video data on the mobile device (iii), the PIP's knowledge is inquired each time a user triggers an event indicating an information flow, e.g., when trying to take a screenshot. The event of taking a screenshot is intercepted by the mobile device's PEP and is only allowed if the application in the foreground does not access the video stream protected by our policy. By this means we implement the requirement [**req08**]. Note that this policy enforcement purposefully protects video data originating from our fall detection system: It does not interfere with the usage of any other video data on the mobile device.

VI. Conclusions

We introduced a fall detection system, which has the potential of improving patient safety even in times that are characterized by a severe shortage of qualified nursing staff. Its transparent design and operation contributes to its acceptance among patients, visitors, and staff members. We also aligned the design with requirements of German data protection law in order to obtain legal certainty. Using cryptographic schemes for authenticating staff members and encrypting video streams as well as deploying distributed usage control combined with information flow tracking capabilities for enforcing anonymization of exposed video data and preventing redistribution of video data, we ensure that the privacy-related requirements of all stakeholders are fulfilled.

References

M. Harvan and A. Pretschner. State-based Usage Control Enforcement with Data Flow Tracking Using System Call Interposition. In 3rd Intl. Conf. on Network and System Security, pages 373 –380, 2009.

M. Hilty, A. Pretschner, D. A. Basin, C. Schaefer, and T. Walter. A Policy Language for Distributed Usage Control. In J. Biskup and J. Lopez, editors, ESORICS, volume 4734 of Lecture Notes in Computer Science, pages 531–546. Springer, 2007.

G. Hornung and M. Desoi. "Smart Cameras" und automatische Verhaltensanalyse. Verfassungs- und datenschutzrechtliche Probleme der nächsten Generation der Videoüberwachung. Kommunikation & Recht, 3:153–158, 2011.

F. Kelbert and A. Pretschner. Data Usage Control Enforcement in Distributed Systems. In Proc. CODASPY, pages 71–82, 2013.

N. Noury, A. Fleury, P. Rumeau, A. Bourke, G. Laighin, V. Rialle, and J. Lundy. Fall Detection - Principles and Methods. In Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International

Conference of the IEEE, pages 1663–1666, Aug 2007.

A. Pretschner, M. Hilty, and D. A. Basin. Distributed Usage Control. Commun. ACM, 49(9):39–44, 2006.

A. Pretschner, E. Lovat, and M. Büchler. Representation-independent Data Usage Control. In Data Privacy Management and Autonomous Spontaneus Security - 6th International Workshop, DPM 2011, and 4th International Workshop, SETOP 2011, Leuven, Belgium, September 15-16, 2011, Revised Selected Papers, pages 122–140, 2011.

A. Roßnagel, M. Desoi, and G. Hornung. Gestufte Kontrolle bei Videoüberwachungsanlagen - ein Drei-Stufen-Modell als Vorschlag zur grundrechtsschonenden Gestaltung. Datenschutz und Datensicherheit, 35(10):694–701, 2011.

Birnstill, P.; Bretthauer, S.; Greiner, S. & Krempel, E. Privacy-preserving Surveillance: an Interdisciplinary Approach International Data Privacy Law, Oxford University Press, 2015