How *Accountability* is Implemented and Understood in Research Tools

A Systematic Mapping Study

Severin Kacianka[1], Kristian Beckers[2], Florian Kelbert[3], and Prachi Kumari

[1]Technical University of Munich, [2]Siemens, [3]Imperial College London
kacianka@in.tum.de, kristian.beckers@siemens.com,
f.kelbert@imperial.ac.uk, prachi.kumari@tum.de

Abstract. [Context/ Background]: With the increasing use of cyber-physical systems in complex socio-technical setups, mechanisms that hold specific entities accountable for safety and security incidents are needed. Although there exist models that try to capture and formalize accountability concepts, many of these lack practical implementations. We hence know little about how accountability mechanisms work in practice and how specific entities could be held responsible for incidents. **[Goal]**: As a step towards the practical implementation of providing accountability, this systematic mapping study investigates existing implementations of accountability concepts with the goal to (1) identify a common definition of accountability and (2) identify the general trend of practical research. **[Method]:** To survey the literature for existing implementations, we conducted a systematic mapping study. **[Results]:** We thus contribute by providing a systematic overview of current accountability realizations and requirements for future accountability approaches. **[Conclusions]:** We find that existing practical accountability research lacks a common definition of accountability in the first place. The research field seems rather scattered with no generally accepted architecture and/or set of requirements. While most accountability implementations focus on privacy and security, no safety-related approaches seem to exist. Furthermore, we did not find excessive references to relevant and related concepts such as reasoning, log analysis and causality.

Keywords: accountability, tools, literature review, survey, systematic mapping study

1 Introduction

Traditionally, IT practitioners have aimed to avoid safety and security incidents using preventive measures. In complex systems, however, it is often hard to enumerate and plan for possible contingencies. Besides, preventive measures generally require many additional resources and are expensive to implement [17]. As a consequence, the focus of research has shifted towards alternative ideas like detective security [23] or root cause analysis [24]. Detective security is inspired

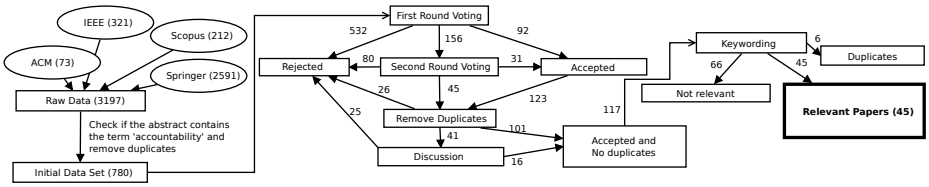

Fig. 1: The "sieving" process

by how law enforcement works in the real world [26]: Speeding violations are not prevented by technical means, e.g. by limiting the maximum speed of the car, but by punishment if caught exceeding the speed limit.

To develop a broad and structured understanding of these and related issues and research undertakings, we designed a mapping study with a focus on *accountability* in the context of privacy, safety, and security. We thus survey the literature of practical accountability implementations that address violations of safety, security, and privacy requirements with the goal to identify the set of existing methods and approaches. Our focus is on the post-mortem analysis of unwanted events.

In terms of related work, Xiao et al. [30] investigate accountability in computer networks and distributed systems. In contrast to their work, we focus on implementations and do not restrict our study to computer networks. While Papanikolaou and Pearson [19] give a cross-discipline overview of the term accountability, they focus on theoretical definitions and do not consider applications.

Our **contribution** is a systematic mapping study on accountability in the context of privacy, safety, and security requirements. We identify which contributions were made over time, the various application domains, layers of abstraction, technologies and protocols in implementing accountability in socio-technical systems. We find that even though there exist very few tools for accountability, it is a growing area of research in different domains. All **raw data** of our study can be found online [16]; see `https://acc.in.tum.de/accountability_2016/` for a more interactive viewer of the data.

2 Methodology

We followed the five-step methodology laid out by Petersen et. al. [21]: (1) definition of research questions (Section 2.1), (2) conduct search (Section 2.2), (3) screening of papers (Section 2.3), (4) keywording using abstracts (Section 2.4), and (5) data extraction and mapping process (Section 2.5). This section describes our instantiation of this methodology. All steps were conducted jointly by the four authors of this paper. The later stages (screening, keywording and mapping) were conducted using a custom written web tool, that offered all authors a unified interface and functioned as a review tool. Figure 1 illustrates our process.

2.1 Definition of Research Questions

We were interested in answering the following research questions:

RQ1 What types of research papers have been published over the years?
RQ2 Which application domains have seen most implementations?
RQ3 Which underlying techniques/protocols are implemented by these tools, at which layers of abstraction are these tools deployed and is there a trend?
RQ4 What do the underlying definitions of accountability have in common?
RQ5 Are prominent contributors recognizable? How are they related to each other?

2.2 Paper Search

In accordance with our research questions, we constructed the search string *accountability AND (privacy OR safety OR security) AND (tool OR implementation OR application)* and adapted it to the idiosyncrasies of each digital library. We limited our search to those technical domains, because we know that accountability is a focus in those fields and because otherwise the result set balloons, encompassing mostly papers covering (non-technical) management and governance problems.

We obtained a basic set of publications from ACM [1] (73 results), IEEE [2] (321), Scopus [3] (212) and Springer [4] (2591), as shown in Table 1, column 'Raw'. As a first step, we stored the search results as CSV files. For this, IEEE and Scopus provided CSV export functionalities, comprising authors, titles, and abstracts. Springer's export functionality did not include abstracts, hence we used a simple script to access the abstracts from the publication's URL. To extract this information from ACM, we used the Zotero tool [5].

Due to the comparatively large amount of results returned by Springer, we performed an initial screening step for all Springer results. We realized that a large amount of those results did not feature the term "accountability" within their abstract. We thus randomly selected 40 publications that did not refer to accountability in their abstract. As it turned out that none of these publications were indeed related to our study subject, we removed all Springer publications that did not feature the term accountability in their abstract. For consistency, we also did this check for the other sources, but had to remove no papers for that reason. Further, Scopus is a meta-search engine that searches, amongst other sources, also the three primary libraries. Scopus thus introduced duplicates. After an additional screening for duplicates and removing them, we obtained the dataset shown in Table 1, column 'Cleanup'. To be consistent in the removal process, we always kept the Scopus version of a paper. Hence for Scopus the number of papers is the same in the columns 'Raw' and 'Cleanup'.

2.3 Screening

We used a custom collaborative web tool to further screen the remaining 780 papers based on their title, keywords and abstract. In this step, we excluded

Table 1: Dataset overview

Source	Raw	Cleanup	Relevant
ACM	73	45	5
IEEE	321	201	25
Scopus	212	212	5
Springer	2591	322	10
Total	**3197**	**780**	**45**

all publications that (i) did not report a tool, implementation or application, (ii) were not related to privacy, safety, or security, or that (iii) reported only an idea, formalism or abstract framework. To ensure consistent decisions from all reviewers, we had frequent meetings. The first meeting was scheduled after every reviewer completed approximately 10 reviews, follow up meetings were held after approximately 50 reviews per authors. The frequency of meeting slowly decreased after the reviewers got more familiar with the screening process.

In practice, our web tool presented each paper randomly to two (out of four) researchers, who then read the abstract and decided whether to include or exclude the paper based on the above criteria. If the researchers' decision was unanimous, the paper was accepted (92 papers) or rejected (532 papers) accordingly. In a second round, all 156 papers with disagreements were presented to two additional researchers. Upon a clear majority of 3-1, the paper was accepted (31 papers) or rejected (80 papers). After this phase, we manually identified and removed 26 more duplicates.

In the following round, the 41 papers that had received a 2-2 draw were discussed in the presence of all researchers and a final verdict was reached. In this phase 25 papers were rejected. Overall, 117 papers proceeded to the next phase of keywording.

2.4 Keywording

In the keywording phase, we classified the remaining 117 papers. For this, we initialized our custom web tool with an intuitive set of keywords agreed upon by discussion among the authors (e.g., security, monitoring, or cloud). These keywords emerged from the authors' experience during the initial screening phase. We also added some keywords under the category of "sanity check" to further exclude irrelevant papers. These keyword-categories were: "No implementation", "Not about accountability", "Full text not available" (was never used) and "I am not sure, I need help". The last category was used if an author was not sure and wanted to discuss the paper with another author. Each paper was then keyworded by one author. Apart from the above initial keywords, each author was able to create new ones on the fly. To ensure a common understanding of the keywords, we again held regular meetings to discuss the keywords.

Despite the previous screening step, 66 papers had to be removed because they (i) did not describe an implementation or (ii) were not about accountability.

[31] Ahmed and Ahamad (2014)	[32] Alexiou et al. (2013)	[33] Ali and Moreau (2013)
[34] Ali et al. (2014)	[35] Ali et al. (2013)	[36] Asokan et al. (2013)
[37] Brzuska et al. (2014)	[38] Cherrueau and Sudholt (2014)	[39] Choi et al. (2005)
[40] Clifton and Fernandez (1988)	[41] Dailianas et al. (2000)	[42] De Oliveira et al. (2013)
[43] Fahl et al. (2014)	[44] Flegel (2002)	[45] Fugkeaw et al. (2007)
[46] Fugkeaw et al. (2009)	[47] Haidar et al.(2010)	[48] Jedrzejczyk et al. (2010)
[49] Kang et al. (2014)	[50] Khalasi et al. (2012)	[51] Ko et al. (2011)
[52] Ko and Will (2014)	[53] Kuacharoen (2012)	[54] Langheinrich (2002)
[55] Wonjun et al. (2009)	[56] Lin and Chang (2009)	[57] Masmoudi et al. (2014)
[58] Michalas and Komninos (2014)	[59] Mivule et al. (2014)	[60] Mortimer and Cook (2010)
[61] Naessens et al. (2005)	[62] Pato et al. (2011)	[63] Pearce et al. (2005)
[64] Pearson et al. (2009)	[65] Popa et al. (2011)	[66] Rubin (1995)
[67] Ruth et al. (2004)	[68] Sriram et al. (2007)	[69] Such et al. (2012)
[70] Such et al. (2013)	[71] Chun et al. (2013)	[72] Kang et al. (2010)
[73] Yang et al. (2010)	[74] Gang et al. (2012)	[75] Zhou et al. (2010)

Fig. 2: All papers part of this study. The full citations can be found online: `https://acc.in.tum.de/accountability_2016/study_papers.pdf`

This is because in the initial screening process we were only deciding on the basis of the papers' titles, abstracts, and provided keywords. Since on this basis it was often not clear whether a paper described an implementation or not, we decided to accept papers if in doubt. After this process, 45 relevant research papers where subject to our study as shown in Table 1, column 'Relevant', and Figure 2.

2.5 Mapping

During the mapping process, our web tool randomly and equally assigned the 45 accepted papers to the four researchers. Each researcher screened the full text, categorized the paper, and gave a short rationale for the categorization. If the paper did not fit into any existing categories, the researcher could create new categories. All of the categories were shared by all researchers in a "tag-cloud" (for example: *Security*, *Efficiency*, or *Health Care*) that was managed by our collaborative web tool. During the process we had several meetings to discuss new categories and unclear publications.

3 Findings

Types of research papers and distribution over the years (RQ1) Our classification of the contributions is based on the classification scheme by Wieringa et al. [28] which was applied to systematic mapping studies by Peterson et al. [21]. We classify the selected papers strictly according to their criteria, which are: Validation Research, Evaluation Research, Solution Proposal, Philosophical Papers, Opinion Papers and Experience Papers. Table 2 maps the selected papers according to these criteria. We realize that all papers focus on solutions and their evaluations. Note that our mapping study focuses on papers that report on techniques that are actually implemented; we excluded meta studies. Hence, we find

Table 2: Paper categorisation into research type facets; grouped by publisher

Category	ACM	IEEE	Springer	Others
Validation Research	[43]			
Evaluation Research	[32, 35, 65, 75]	[31, 34, 38, 46, 51, 52, 55, 57, 71, 73]	[37, 39, 44, 61]	[70]
Solution Proposal	[48, 50]	[33, 40–42, 45, 47, 49, 58, 60, 62, 64, 66, 68, 72, 74]	[36, 53, 54, 56, 63, 67, 69]	[59]

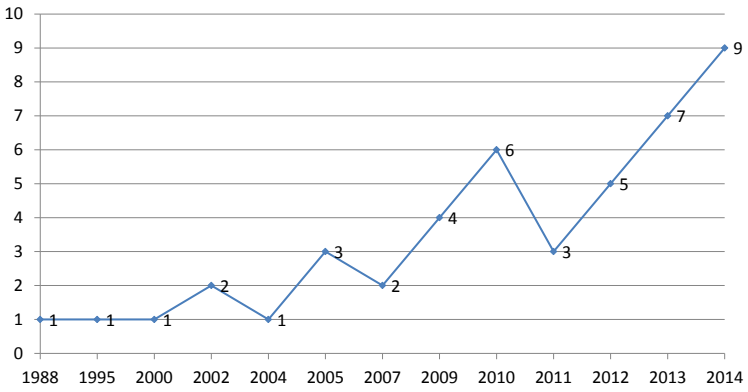

Fig. 3: Number of papers over the years

no papers in the categories experience paper, opinion paper, or philosophical paper.

To identify how the number of contributions developed over time, we analyzed the papers according to their year of publication. Figure 3 shows the graph of the distribution from 1988 to 2014, revealing that accountability implementations started gaining interest in 1988 beginning with the work of [40]. For the first few years until the year 2000, this area did not attract much attention with only three papers in 12 years. There are several crests and troughs starting in the year 2000, but the overall interest of the research community has been increasing. In fact, as shown in Figure 3, every trough is at a higher level than the previous one. Since 2011, there has been a consistent growth in the number of implementations. It is also notable that after the publication of the influential paper by Weitzner et al. [26] (which, as a theoretical paper, is not subject of our study) in 2008, we see relevant publications in every consecutive year.

Interpretation: The research types in the field of implemented accountability approaches are validation, evaluation and solution approaches. It is no surprise that the field started with solution approaches and moved over time to evaluation approaches. The majority of publications in the years 2013 and 2014 are of that type. We have seen only one validation approach. We assume over time the focus of research will go towards evaluation approaches and ultimately valida-

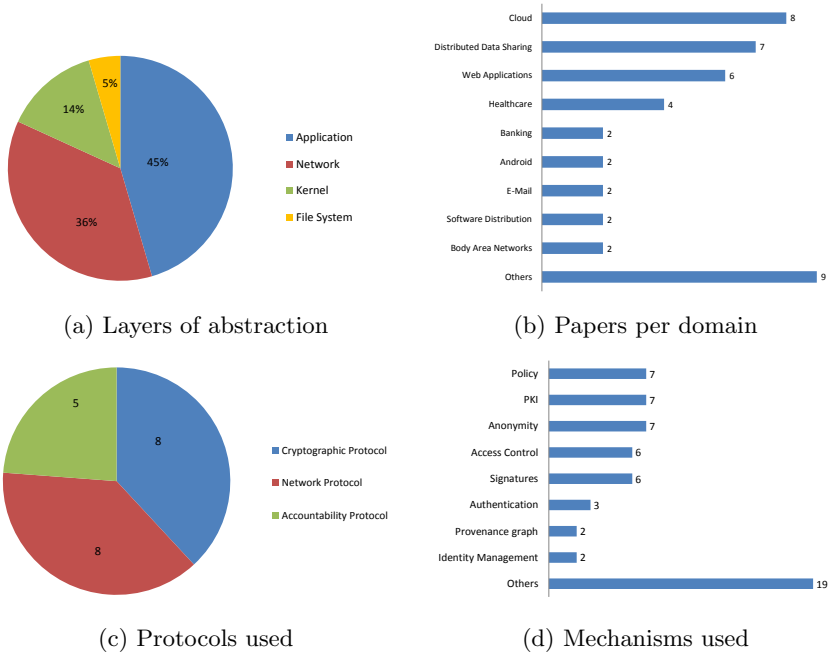

(a) Layers of abstraction

(b) Papers per domain

(c) Protocols used

(d) Mechanisms used

Fig. 4: Findings

tion approaches. Hence, the field evolves towards evaluation research, while we see a clear gap in validation research. Though the initial work on implementing accountability is by [40] in the year 1988, the field of accountability implementations started growing only from the year 2000, as shown in Figure 3. In summary, contributions over the years indicate that accountability is (1) not yet a mature field as indicated by the low number of tools and implementations, and (2) a growing field of research with a consistent increase in the number of tools over the last decade.

Application Domains (RQ2) To answer the second research question, we classified the papers according to the targeted application domains. As shown in Figure 4b, accountability concepts have been mostly implemented for the *cloud* domain with 8 implementations [33, 42, 50–52, 57, 60, 71]. Other important domains are *distributed data sharing* (7 implementations [45,52,59,60,63,67,68]), *web applications* (6; [38, 47, 49, 58, 72, 74]), and *health care* (4; [31, 34, 35, 59]). For other domains we found at most two implementations.

Since the implementation of accountability mechanisms is a relatively new area of research, there are many domains for which only single implementations exist. These have been grouped as *Others* in Figure 4b and include web services, ubiquitous computing, wireless networks, business organization, ecommerce, lottery, insurance, grid computing and location based services.

Interpretation: Cloud computing is en vogue. At the same time, it is one of the application domains where most privacy and data protection concerns have been raised. Distributed data sharing is another such domain. Encryption and access control have been shown to be insufficient for addressing these issues in remote computing and data sharing in general [26]. Hence, it is only obvious that researchers are trying to address privacy and security issues by detective enforcement viz. implementing accountability in these domains. An interesting finding is that web applications and health care domains have not attracted equal focus, especially health care where HIPPA (Health Insurance Portability and Accountability Act of 1996) explicitly mandates accountability enforcements.

Underlying Techniques & Protocols (RQ3) As depicted in Figure 4c, we found three different kinds of protocols that are leveraged by implementations to achieve accountability.

Eight papers use *network protocols* [39,41,42,58,72–75] or *cryptographic protocols* [32, 34, 37, 39, 58, 63, 68, 74], while five papers make use of *accountability protocols* [35,36,61,63,74]. Contrary to our expectation, data provenance protocols are not commonly used for accountability implementations.

Since accountability is the focus of this study, we took a more detailed look at the accountability protocols: [35] uses fingerprinting of wireless connection in body area networks to later proof communication between two parties. [36] describes a system for friends to share resources and uses accountability to prevent abuse. They use internet connection sharing as a use case. [61] treats a similar problem, considering an anonymous e-mail service and providing accountability in case a user abuses the system to commit criminal acts. [63] describes a protocol to resolve disputes about transactions in e-commerce systems. [74] proposes the term "accountable anonymity" and uses an encryption scheme to build an accountable and anonymous internet proxy.

Furthermore, we investigated which mechanisms and techniques are used to implement accountability. As detailed in Figure 4d, we found that most solutions are concerned with *enforcement of policies* (7 solutions [33,38,46,50,54,55,62]), *public key encryption schemes* (7; [32,37,43,46,47,63,66]), *anonymity* (7; [32,44, 58,61,63,65,74]), *access control* (6; [33,36,45,46,55,62]), and *digital signatures* (6; [34,37,42,43,67,68]). Some tools also use *authentication* (3; [45–47]), *provenance graphs* (2; [67,71]), and *identity management* (2; [55,69]) to hold entities accountable in systems. 19 further mechanisms appeared in only one implementation each. These are represented as "Others" in Figure 4d and include certificates, traces, pseudonyms, pseudonymity, log tamper resistance, time synchronization, reputation systems, unlinkability, accountable anonymity, OLAP, questionnaire and report generation, key management, resource description framework (RDF), job-flow tracking, fault detection, monitoring, onion routing, decentralization, and Shamir's threshold scheme.

We further found that accountability mechanisms are mainly implemented at the *application layer* (10 instances; [37,46,50,59–62,65,69,72]) and the *network layer* (8; [34,39,47,52,58,63,72,74]), see Figure 4a. Few solutions are implemented at the *kernel layer* (3; [40,52,71]) and the *file system layer* (1; [44]).

Interpretation: The underlying techniques in accountability implementations are dominated by cryptographic protocols and network protocols. We found only one implementation relying on data provenance and very few accountability-centric protocols which combine, e.g., anonymity with accountability. In addition, we observed three overall trends in mechanisms offered within accountability implementations. First, cryptography is dominating the field with, e.g., public key infrastructures, signature-based solutions, and certificates. Second, access control mechanisms are wide-spread. Either under the term access control or in supporting topics such as policy-based approaches, authentication mechanisms, or identity management. Third, privacy is a recurring theme in particular with respect to anonymity. Further privacy goals such as pseudonymity and unlinkability are supported as well, but to a lesser extend. We sparsely encountered further supporting mechanisms such as provenance and traceability.

Definitions of Accountability (RQ4) We scanned all 45 papers for the definition of accountability. To find the definition, we searched the documents for all occurrences of the term "accountability". We then read the text before and after the highlighted term and looked for a definition.

We found that 20 of the 45 papers provide no explicit definition of accountability. 17 papers provide their own definition, not taking other sources into account. These definitions define accountability in terms of responsibility/ assigning blame (6), non-repudiation/ integrity (3), a-posteriori enforcement (3), collect evidence (2), transparency (2), traceability (1).

Only 8 papers rely on a previously-published and peer-reviewed definition:

Anderson et al. [6] the "(...) ability to associate an action with the responsible entity"

Bhargav-Spantzel et al. [7] "(...) the ability of holding entities responsible for their actions"

Brzuska et al. [8] "A sanitizable signature scheme satisfies non-interactive public accountability, if and only if for a valid message/signature pair (m, σ), a third party can correctly decide whether (m, σ) originates from the signer or from the sanitizer without interacting with the signer or sanitizer."

Ko et al. [18] who rely on [20] and use the definition from the "The Best Practices Act of 2010" (we, however, could not find the formulation in the original source): "the obligation and/ or willingness to demonstrate and take responsibility for performance in light of agreed-upon expectations."

Pearson [20] relies on Weitzner et al. [26] and extends the definition of the "Galway project": "Accountability is the obligation to act as a responsible steward of the personal information of others, to take responsibility for the protection and appropriate use of that information beyond mere legal requirements, and to be accountable for any misuse of that information."

Xiao [29] "Accountability implies that any entity should be held responsible for its own specific action or behavior so that the entity is part of larger chains of accountability. One of the goals of accountability is that once an event has transpired, the events that took place are traceable so that the causes can be determined afterward."

Fig. 5: Collaboration map. The size of nodes and author names corresponds with the author's number of papers (1–3) considered in this study.

These definitions, like the 17 definitions provided by the other papers, are not peer-reviewed and rely on a common understanding of the (dictionary-)meaning[1] of accountability.

Interpretation: It was surprising that no clear and accepted definition of accountability emerged. We assume that the main reason for this is that it is a common English word and everyone has some intuitive understanding of the term. The lack of a clear definition and differentiation from other terms like "responsibility" or "detection" hinders the scientific discourse and the comparability of the approaches. We hope that in the future works will rely on a peer reviewed definition of accountability and that thus trends and relations among approaches will become more pronounced. Despite this, all definitions see accountability as some form of a-posteriori mechanism to provide evidence and ultimately assign blame or responsibility. It relies either on logs or some other form of monitor.

Contributors and Relationships (RQ5)

Collaboration Networks. We analyzed the author networks of the selected papers. First, we find that most authors feature only one publication on ac-

[1] The Oxford dictionary defines accountability as "The fact or condition of being accountable; responsibility". For a more detailed discussion see [19].

Table 3: Most influential researchers.

Name	**Institution**	**Cit.**
Siani Pearson	HP Labs Bristol, UK	16
David L. Chaum	Voting Systems Institute	14
Margo Seltzer	Harvard University, Cambridge, MA, USA	13
Jan Camenisch	IBM Research, Zurich, Switzerland	13
Markus Kirchberg	National University of Singapore, Singapore	11
Kiran Kumar Muniswamy-Reddy	Harvard University, Cambridge, MA, USA	9
Lorrie Faith Cranor	Carnegie Mellon University, Pittsburgh, PA, USA	9
Elisa Bertino	Purdue University, West Lafayette, Indiana, USA	8
Uri J. Braun	Harvard University, Cambridge, MA, USA	8
Gene Tsudik	University of California, Irvine, California, USA	8
Anna Lysyanskaya	Brown University, Providence, RI, USA	8
Wade Trappe	Rutgers University, Piscataway, New Jersey, USA	7
Ian T. Foster	University of Chicago, Chicago, IL, USA	7
Peter Macko	Harvard University, Cambridge, MA, USA	7
Susan Hohenberger	Johns Hopkins University, Baltimore, MD, USA	7

countability implementations, as indicated by the size of the nodes in Figure 5. 13 authors feature two publications, while only one author features three. For the authors with at least two publications, we found that the corresponding papers are closely related follow-up papers. As also indicated by Figure 5, the analyzed author network is very scattered. The authors of accountability tools do not collaborate across research groups. Again, the only papers published by the same authors are [34, 35], [45, 46], and [51, 52, 71] all of which are a series of papers.

These results lead us to the conclusion and hypothesis that the field of accountability implementations would greatly benefit from more systematic collaborations and research among the identified researchers.

Most Influential Researchers. We further analyzed the references of the 45 selected papers. Our goal was to find out whether they share common literature that is essential for the understanding and implementation of accountability mechanisms. Because some authors made heavy use of self citations, we decided to exclude any self references. We realized that there exist some researchers that are cited across many of the study papers. Table 3 shows those researchers that were cited at least seven times.

Interpretation: In contrast to the theoretical discussions of accountability, where we often find citations to papers like the one by Weitzner et al. [26] or Feigenbaum et al. [12], there are no especially noticeable contributors. We assume that there are more prominent works on topics related to (but not called) accountability, like fault localization or root cause analysis. This suggests that a clear and thorough overview of the whole field of computer science is needed. This should then yield to a clearer definition and taxonomy of the term accountability and its related concepts.

4 Synthesis

4.1 Definition of Accountability

One of the main motivation for us to conduct this mapping study was to come to a unified definition of accountability. We originally anticipated that most papers would agree on a specific definition; we assumed it would be the definition of *Information Accountability* as given by Weitzner et al. [26]. We did not expect that most papers would use the term without any definition or that so many papers would use ad-hoc definitions. Yet, this diversity of definitions also highlights the different facets of accountability and can serve as a basis for a more general definition. Analyzing all given definitions, we can identify five main themes:

1. Accountability should associate (or *link*) actions to entities (often individuals).
2. This link should then be used (often by a neutral third party) to hold the entity *responsible* for that action (often the terms blame and punish are used).
3. All definitions implicitly rely on some notion of *log* that is complete, tamper-proof and available to the neutral third party.
4. Another implicit assumption is that the log data can be used to *reason* about the events that have transpired.
5. All definitions only consider *single* systems. There is no notion of "distributed" accountability in those definitions.

Considering these aspects, we propose the following work-in-progress definition of accountability:

1. *Accountability* is a property of a system or a collection of systems and is ensured by an *Accountability Mechanism*.
2. An *Accountability Mechanism* is part of an *Accountable System* and reasons over a tamper-proof log to link effects of that system to entities.
3. An entity is (partially) *accountable* for a given effect if an *Accountability Mechanism* can prove a causal link between the entity's action and the given effect.
4. The set of entities *accountable* for a given effect is the set of all entities for which an *Accountability Mechanism* can prove a causal link between the entities' actions and the given effect.

4.2 Future Research Directions

We identified two main observations from the 45 study papers:

1. Preventing unwanted behavior is increasingly difficult in distributed and highly interconnected systems.
2. The impact of any unwanted behavior of computer systems increases with their adoption.

The first observation is corroborated by the domains that accountability mechanisms are mostly used in: cloud computing, distributed data sharing and web applications are all highly distributed systems. The use for accountability in a single user system is limited: as long as the system is not faulty, any effect is the result of its sole user's actions. Consequently, we expect a rising demand for accountability and its implementations in the fields of cyber-physical systems, smart systems, and similar fields where devices are only now being connected to form a wider Internet of Things. Indeed, a recent position paper by Datta et al. [11] calls for exactly such mechanisms to enhance the security of cyber-physical systems.

The second observation is best illustrated with the surprisingly high number of papers from the health care and medical domain. In our opinion, this can be explained with the legal risks and liabilities within the field. Medical devices are highly regulated and malfunctioning can be a serious threat to life and limb. If a pacemaker malfunctions, it is impossible to simply reboot the system or to restore the last backup. Similarly, computer systems already control cars, drones and hydro-dams. Any malfunctioning can have serious consequences and thus a high risk of legal action. In such a case the operator (and often also regulatory bodies) want a clear trace of accountability.

5 Threats to Validity

There are three main threats to validity of this mapping study: the selection of papers, our potential bias when reviewing and categorizing the papers, and the timeliness of the data.

Selection of papers. By limiting our study to the term "accountability", we might have missed papers that implement similar concepts but refer to them by different terms (e.g., "black box" or "root cause analysis"). We made our choice based on experiences of existing research. Petticrew and Roberts [22] highlight that the two main issues in conducting a literature survey are the sensitivity and specificity of the search. The sensitivity refers to the number of relevant publications of a search. Specificity describes the number of irrelevant studies of a search. The aim is to have a high sensitivity and a low specificity of a search. Synonyms may increase the sensitivity, but it also increases the specificity. Previous experiences of literature studies advocate simple search strings and limited synonyms to achieve an optimal trade-off between specificity and sensitivity [25].

Potential bias. It is possible that we collectively misclassified some papers. We countered this with a multi-staged voting process and took special care that every paper was reviewed by at least two different researchers. Furthermore, an inherent limitation of mapping studies is the superficial review of the source literature. Especially in the early stages we only looked at the abstract of a paper and not at its content. In the later stages, however, we examined each paper more carefully.

Timeliness of data. A well-known problem with literature reviews is that they are quickly outdated. The present data was gathered in 2015 and contains

works up to the year 2015. This means that any more recent works are not part of our dataset. A recent (June 2017) manual check of the publishers' digital databases with study's search string returned one additional survey about accounting in content distribution networks [10] and some additional implementations in the field of e-health [13,14,27] and cloud computing [9,15]. While this search was not backed by a systematic process, we have not found any indication that our study's conclusions need revision. On the contrary, this cursory search seems to confirm our findings.

6 Conclusion & Future Work

Through this systematic mapping study, we establish the state of the art in accountability implementations and tools.

We have considered only those papers that describe an implementation. We did not consider contributions that described, even if in detail, how the ideas *could* be implemented. In this context, an interesting finding is that none of the papers have evaluated their tools for performance. This is important because one key factor that could limit the usefulness of accountability mechanisms is performance efficiency. The reason is that the origin of unwanted events is typically tracked using logging and analysis of "interesting" system events. Depending on the complexity of the analysis algorithm and the size of the logs, accountability implementations could be very expensive in terms of computation. It would help to get an insight into how the existing implementations perform and if the concepts can be reused in domains where real-time processing is needed, e.g., the automotive domain.

Another identified gap is the missing link between the high-level unwanted events that take place in an environment (e.g., personal and medical data is leaked in a Healthcare domain application) and the low-level unwanted events that are logged in the running technical systems (e.g., system calls reading from confidential files and writing to a socket in a network connection). It is important to establish this link because unwanted events are extracted from high-level requirements of privacy, security and safety properties and there is no universally agreed upon semantics of the relevant high-level events (e.g., data leak) in terms of low-level technical events (e.g., system calls writing to sockets). Though this gap has been filled in the context of preventive enforcement of usage control, it is not clear how this could be done for accountability.

One of our goals of this study was to identify which properties are often considered in combination with accountability. We found that security and privacy are most often considered along with accountability. Other important properties are integrity, provenance, trust, legal compliance, confidentiality, transparency, traceability, auditability and non-repudiation. While most papers consider more than one of these properties, an interesting finding is that none of the papers implement a safety property. This discovery points out a gap in the work on accountability for safety-critical systems.

We were also surprised that relevant concepts like reasoning, log analysis and causality did not feature prominently in the result set. Current accountability technologies focus mainly on preventive concepts (policies and access control) or authenticity/Non-repudiation (public key infrastructures, anonymity and signatures). At the high-level view of this mapping study we could not reliably identify an a-posteriori approach. We believe that this needs to change in the future: While it is feasible to manually analyze the logs (flight recorders) the few times a year an aircraft crashes, it becomes infeasible when multiple drones crash every day.

Our conclusion is that though accountability concepts have been around for quite some time, this area has not seen enough implementations, especially of a-posteriori approaches. At the technical level, there exists no generally accepted architecture and we did not come across contributions that give insights into acceptability issues like usability, scalability, etc. At the methodological level, there are no processes for deriving accountability-specific requirements. Thus, there is plenty of room for developing accountability infrastructures.

Acknowledgments. This work was funded in part by the Munich Center for Internet Research and the TUM Living Lab Connected Mobility (TUM LLCM) project which has been funded by the Bavarian Ministry of Economic Affairs and Media, Energy and Technology (StMWi) through the Center Digitisation.Bavaria, an initiative of the Bavarian State Government.

References

1. Acm digital library. `http://dl.acm.org/` (2017). [Online; accessed 2017-06-07]
2. IEEE Xplore. `http://ieeexplore.ieee.org` (2017). [Online; accessed 2017-06-07]
3. Scopus. `http://www.scopus.com` (2017). [Online; accessed 2017-06-07]
4. Springer. `http://link.springer.com` (2017). [Online; accessed 2017-06-07]
5. Zotero. `http://www.zotero.org` (2017). [Online; accessed 2017-06-07]
6. Andersen, D.G., Balakrishnan, H., Feamster, N., Koponen, T., Moon, D., Shenker, S.: Accountable internet protocol (aip). In: ACM Computer Communication Review, vol. 38, pp. 339–350. ACM (2008)
7. Bhargav-Spantzel, A., Camenisch, J., Gross, T., Sommer, D.: User centricity: a taxonomy and open issues. Journal of Computer Security **15**(5), 493–527 (2007)
8. Brzuska, C., Pöhls, H.C., Samelin, K.: Non-interactive public accountability for sanitizable signatures. In: Public Key Infrastructures, Services and Applications, pp. 178–193. Springer (2012)
9. Chen, H., Tu, S., Zhao, C., Huang, Y.: Provenance cloud security auditing system based on log analysis. In: 2016 IEEE International Conference of Online Analysis and Computing Science (ICOACS), pp. 155–159 (2016). DOI 10.1109/ICOACS.2016.7563069
10. Coileáin, D.O., O'mahony, D.: Accounting and accountability in content distribution architectures: A survey. ACM Comput. Surv. **47**(4), 59:1–59:35 (2015). DOI 10.1145/2723701. URL `http://doi.acm.org/10.1145/2723701`
11. Datta, A., Kar, S., Sinopoli, B., Weerakkody, S.: Accountability in cyber-physical systems. In: 2016 Science of Security for Cyber-Physical Systems Workshop (SOSCYPS), pp. 1–3 (2016). DOI 10.1109/SOSCYPS.2016.7579998

12. Feigenbaum, J., Jaggard, A.D., Wright, R.N.: Towards a formal model of accountability. In: Workshop on new security paradigms workshop, pp. 45–56. ACM (2011)
13. Grunwel, D., Sahama, T.: Delegation of access in an information accountability framework for ehealth. In: Proceedings of the Australasian Computer Science Week Multiconference, ACSW '16, pp. 59:1–59:8. ACM, New York, NY, USA (2016). DOI 10.1145/2843043.2843383
14. Grunwell, D., Batista, P., Campos, S., Sahama, T.: Managing and sharing health data through information accountability protocols. In: 2015 17th International Conference on E-health Networking, Application Services (HealthCom), pp. 200–204 (2015). DOI 10.1109/HealthCom.2015.7454498
15. Jain, J.R., Asaduzzaman, A.: A novel data logging framework to enhance security of cloud computing. In: SoutheastCon 2016, pp. 1–6 (2016). DOI 10.1109/SECON.2016.7506764
16. Kacianka, S., Beckers, K., Kelbert, F., Kumari, P.: Dataset: How Accountability is Understood and Realized (2017). DOI 10.5281/zenodo.807129. URL `https://doi.org/10.5281/zenodo.807129`
17. Kelbert, F., Pretschner, A.: A Fully Decentralized Data Usage Control Enforcement Infrastructure. In: Applied Cryptography and Network Security, *Lecture Notes in Computer Science*, vol. 9092, pp. 409–430. Springer International Publishing (2015)
18. Ko, R.K., Jagadpramana, P., Mowbray, M., Pearson, S., Kirchberg, M., Liang, Q., Lee, B.S.: Trustcloud: A framework for accountability and trust in cloud computing. In: IEEE World Congress on Services, pp. 584–588. IEEE (2011)
19. Papanikolaou, N., Pearson, S.: A cross-disciplinary review of the concept of accountability. In: Proceedings of the International Workshop on Trustworthiness, Accountability and Forensics in the Cloud (TAFC) (2011)
20. Pearson, S.: Toward accountability in the cloud. IEEE Internet Computing **15**(4), 64 (2011)
21. Petersen, K., Feldt, R., Mujtaba, S., Mattsson, M.: Systematic mapping studies in software engineering. In: 12th Intl. Conf. on Evaluation and Assessment in Software Engineering, vol. 17. sn (2008)
22. Petticrew, M., Roberts, H.: Systematic Review in the Social Sciences: A Practical Guide. Blackwell Publishing (2006)
23. Povey, D.: Optimistic security: A new access control paradigm. In: Proceedings of the 1999 Workshop on New Security Paradigms, pp. 40–45. ACM (2000)
24. Rooney, J.J., Heuvel, L.N.V.: Root cause analysis for beginners. Quality progress **37**(7), 45–56 (2004)
25. Salleh, N., Mendes, E., Grundy, J.: Empirical studies of pair programming for cs/se teaching in higher education: A systematic literature review. IEEE Transactions on Software Engineering **37**(4), 509–525 (2011)
26. Weitzner, D.J., Abelson, H., Berners-Lee, T., Feigenbaum, J., Hendler, J., Sussman, G.J.: Information accountability. Commun. ACM **51**(6), 82–87 (2008)
27. Wickramage, C., Sahama, T., Fidge, C.: Anatomy of log files: Implications for information accountability measures. In: Healthcom, pp. 1–6 (2016). DOI 10.1109/HealthCom.2016.7749426
28. Wieringa, R., Maiden, N., Mead, N., Rolland, C.: Requirements engineering paper classification and evaluation criteria: A proposal and a discussion. Requir. Eng. **11**(1), 102–107 (2005)
29. Xiao, Y.: Flow-net methodology for accountability in wireless networks. Network, IEEE **23**(5), 30–37 (2009)
30. Xiao, Z., Kathiresshan, N., Xiao, Y.: A survey of accountability in computer networks and distributed systems. Security and Communication Networks (2012)

Study Papers

31. Ahmed, M., Ahamad, M.: Combating Abuse of Health Data in the Age of eHealth Exchange. In: IEEE Int. Conf. on Healthcare Informatics, pp. 109–118 (2014)
32. Alexiou, N., Laganà, M., Gisdakis, S., Khodaei, M., Papadimitratos, P.: VeSPA: Vehicular Security and Privacy-preserving Architecture. In: 2nd ACM Workshop on Hot Topics on Wireless Network Security and Privacy, pp. 19–24. ACM (2013)
33. Ali, M., Moreau, L.: A Provenance-Aware Policy Language (cProvl) and a Data Traceability Model (cProv) for the Cloud. In: Third International Conf. on Cloud and Green Computing, pp. 479–486 (2013)
34. Ali, S., Sivaraman, V., Ostry, D., Tsudik, G., Jha, S.: Securing First-Hop Data Provenance for Bodyworn Devices Using Wireless Link Fingerprints. IEEE Transactions on Information Forensics and Security **9**(12), 2193–2204 (2014)
35. Ali, S.T., Sivaraman, V., Ostry, D., Jha, S.: Securing Data Provenance in Body Area Networks Using Lightweight Wireless Link Fingerprints. In: Proc. 3rd International Workshop on Trustworthy Embedded Devices, pp. 65–72. ACM (2013)
36. Asokan, N., Dmitrienko, A., Nagy, M., Reshetova, E., Sadeghi, A.R., Schneider, T., Stelle, S.: CrowdShare: Secure Mobile Resource Sharing. In: Applied Cryptography and Network Security, *LNCS*, vol. 7954, pp. 432–440. Springer (2013)
37. Brzuska, C., Pöhls, H., Samelin, K.: Efficient and Perfectly Unlinkable Sanitizable Signatures without Group Signatures. In: Public Key Infrastructures, Services and Applications, *LNCS*, vol. 8341, pp. 12–30. Springer Berlin Heidelberg (2014)
38. Cherrueau, R.A., Sudholt, M.: Enforcing Expressive Accountability Policies. In: IEEE 23rd International WETICE Conference, pp. 333–338 (2014)
39. Choi, C., Dong, Y., Zhang, Z.L.: LIPS: Lightweight Internet Permit System for Stopping Unwanted Packets. In: NETWORKING 2005. Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communications Systems, *LNCS*, vol. 3462, pp. 178–190. Springer (2005)
40. Clifton, D., Fernandez, E.: A Microprocessor Design for Multilevel Security. In: Fourth Aerospace Computer Security Applications Conference, pp. 194–198 (1988)
41. Dailianas, A., Yemini, Y., Florissi, D., Huang, H.: MarketNet: market-based protection of network systems and services-an application to SNMP protection. In: Proc. 19th Annual Joint Conference of the IEEE Computer and Communications Societies., vol. 3 (2000)
42. De Oliveira, A., Sendor, J., Garaga, A., Jenatton, K.: Monitoring Personal Data Transfers in the Cloud. In: IEEE 5th Intl. Conf. on Cloud Computing Technology and Science, vol. 1, pp. 347–354 (2013)
43. Fahl, S., Dechand, S., Perl, H., Fischer, F., Smrcek, J., Smith, M.: Hey, NSA: Stay Away from My Market! Future Proofing App Markets Against Powerful Attackers. In: Proc. 2014 ACM Conference on Computer and Communications Security, pp. 1143–1155. ACM (2014)
44. Flegel, U.: Pseudonymizing Unix Log Files. In: Infrastructure Security, *LNCS*, vol. 2437, pp. 162–179. Springer (2002)
45. Fugkeaw, S., Manpanpanich, P., Juntapremjitt, S.: AmTRUE: Authentication Management and Trusted Role-based Authorization in Multi-Application and Multi-User Environment. In: The International Conf. on Emerging Security Information, Systems, and Technologies, pp. 216–221 (2007)
46. Fugkeaw, S., Manpanpanich, P., Juntapremjitt, S.: A-COLD: Access Control of Web OLAP over Multi-data Warehouse. In: International Conf. on Availability, Reliability and Security, pp. 469–474 (2009)

47. Haidar, A., Zasada, S., Coveney, P., Abdallah, A., Beckles, B.: Audited Credential Delegation - A User-centric Identity Management Solution for Computational Grid Environments. In: Sixth International Conf. on Information Assurance and Security, pp. 222–227 (2010)
48. Jedrzejczyk, L., Price, B.A., Bandara, A.K., Nuseibeh, B.: On the Impact of Real-time Feedback on Users' Behaviour in Mobile Location-sharing Applications. In: Proc. Sixth Symposium on Usable Privacy and Security, pp. 14:1–14:12. ACM (2010)
49. Kang, Y., Schiffman, A., Shrager, J.: RAPPD: A Language and Prototype for Recipient-Accountable Private Personal Data. In: IEEE Security and Privacy Workshops, pp. 49–56 (2014)
50. Khalasi, G., Chaudhari, M.: TrustGK Monitor: 'Customer Trust As a Service' for the Cloud. In: Proc. CUBE International Information Technology Conference, pp. 537–543. ACM (2012)
51. Ko, R., Jagadpramana, P., Lee, B.S.: Flogger: A File-Centric Logger for Monitoring File Access and Transfers within Cloud Computing Environments. In: IEEE 10th International Conf. on Trust, Security and Privacy in Computing and Communications, pp. 765–771 (2011)
52. Ko, R., Will, M.: Progger: An Efficient, Tamper-Evident Kernel-Space Logger for Cloud Data Provenance Tracking. In: IEEE 7th International Conf. on Cloud Computing, pp. 881–889 (2014)
53. Kuacharoen, P.: Design and Implementation of a Secure Online Lottery System. In: Advances in Information Technology, *Communications in Computer and Information Science*, vol. 344, pp. 94–105. Springer (2012)
54. Langheinrich, M.: A Privacy Awareness System for Ubiquitous Computing Environments. In: UbiComp 2002: Ubiquitous Computing, *LNCS*, vol. 2498, pp. 237–245. Springer (2002)
55. Lee, W., Squicciarini, A., Bertino, E.: The Design and Evaluation of Accountable Grid Computing System. In: 29th IEEE International Conference on Distributed Computing Systems, pp. 145–154 (2009)
56. Lin, K.J., Chang, S.: A Service Accountability Framework for QoS Service Management and Engineering. Information Systems and e-Business Management **7**(4), 429–446 (2009)
57. Masmoudi, F., Loulou, M., Kacem, A.: Multi-tenant Services Monitoring for Accountability in Cloud Computing. In: IEEE 6th Intl. Conf. on Cloud Computing Technology and Science, pp. 620–625 (2014)
58. Michalas, A., Komninos, N.: The Lord of the Sense: A Privacy Preserving Reputation System for Participatory Sensing Applications. In: IEEE Symp. on Computers and Communication, pp. 1–6 (2014)
59. Mivule, K., Otunba, S., Tripathy, T.: Implementation of Data Privacy and Security in an Online Student Health Records System. Tech. rep., Department of Computer Science, Bowie State University (2014)
60. Mortimer, D., Cook, N.: Supporting Accountable Business to Business Document Exchange in the Cloud. In: IEEE International Conf. on Service-Oriented Computing and Applications, pp. 1–8 (2010)
61. Naessens, V., De Decker, B., Demuynck, L.: Accountable Anonymous E-Mail. In: Security and Privacy in the Age of Ubiquitous Computing, *IFIP Advances in Information and Communication Technology*, vol. 181, pp. 3–18. Springer (2005)
62. Pato, J., Paradesi, S., Jacobi, I., Shih, F., Wang, S.: Aintno: Demonstration of Information Accountability on the Web. In: IEEE 3rd Intl. Conf. on Privacy,

Security, Risk and Trust and 2011 IEEE 3rd Intl. Conf. on Social Computing, pp. 1072–1080 (2011)
63. Pearce, C., Bertok, P., Van Schyndel, R.: Protecting Consumer Data in Composite Web Services. In: Security and Privacy in the Age of Ubiquitous Computing, *IFIP AICT*, vol. 181, pp. 19–34. Springer US (2005)
64. Pearson, S., Rao, P., Sander, T., Parry, A., Paull, A., Patruni, S., Dandamudi-Ratnakar, V., Sharma, P.: Scalable, accountable privacy management for large organizations. In: 13th Enterprise Distributed Object Computing Conference Workshops, pp. 168–175 (2009)
65. Popa, R.A., Blumberg, A.J., Balakrishnan, H., Li, F.H.: Privacy and Accountability for Location-based Aggregate Statistics. In: Proc. 18th ACM Conf. on Computer and Communications Security, pp. 653–666. ACM (2011)
66. Rubin, A.: Trusted Distribution of software over the Internet. In: Proc. Symp. on Network and Distributed System Security, pp. 47–53 (1995)
67. Ruth, P., Xu, D., Bhargava, B., Regnier, F.: E-notebook Middleware for Accountability and Reputation Based Trust in Distributed Data Sharing Communities. In: Trust Management, *LNCS*, vol. 2995, pp. 161–175. Springer (2004)
68. Sriram, V., Narayan, G., Gopinath, K.: SAFIUS - A Secure and Accountable Filesystem over Untrusted Storage. In: Fourth International IEEE Security in Storage Workshop, pp. 34–45 (2007)
69. Such, J.M., Espinosa, A., Garcia-Fornes, A.: An Agent Infrastructure for Privacy-Enhancing Agent-Based E-commerce Applications. In: Advanced Agent Technology, *LNCS*, vol. 7068, pp. 411–425. Springer Berlin Heidelberg (2012)
70. Such, J.M., García-Fornes, A., Espinosa, A., Bellver, J.: Magentix2: A privacy-enhancing Agent Platform. Engineering Applications of Artificial Intelligence **26**(1), 96–109 (2013)
71. Suen, C.H., Ko, R., Tan, Y.S., Jagadpramana, P., Lee, B.S.: S2Logger: End-to-End Data Tracking Mechanism for Cloud Data Provenance. In: 12th IEEE International Conf. on Trust, Security and Privacy in Computing and Communications, pp. 594–602 (2013)
72. Wang, K., Malozemoff, A., Jia, N., Han, C., Maheswaran, M.: A Social Accountability Framework for Computer Networks. In: IEEE Global Telecommunications Conference, pp. 1–6 (2010)
73. Xiao, Y., Meng, K., Takahashi, D.: Implementation and Evaluation of Accountability Using Flow-net in Wireless Networks. In: Military Communications Conference, pp. 7–12 (2010)
74. Xu, G., Aguilera, L., Guan, Y.: Accountable Anonymity: A Proxy Re-Encryption Based Anonymous Communication System. In: IEEE 18th Intl. Conf. on Parallel and Distributed Systems, pp. 109–116 (2012)
75. Zhou, W., Sherr, M., Tao, T., Li, X., Loo, B.T., Mao, Y.: Efficient Querying and Maintenance of Network Provenance at Internet-scale. In: Proc. 2010 ACM SIGMOD International Conf. on Management of Data, pp. 615–626. ACM (2010)