DAGSTUHL
MANIFESTOS

Volume 1, Issue 1, January – December 2011

ISSN 2193-2433

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany.
Online available at http://www.dagstuhl.de/dagman

Publication date
March, 2012

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at `http://dnb.d-nb.de`.

Aims and Scope
The manifestos from Dagstuhl Perspectives Workshops are published in the *Dagstuhl Manifestos* journal. Each manifesto aims for describing the state-of-the-art in a field along with its shortcomings and strenghts. Based on this, position statements and perspectives for the future are illustrated. A manifesto typically has a less technical character; instead it provides guidelines and roadmaps for a sustainable organisation of future progress.

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Manifestos, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
publishing@dagstuhl.de

Manifesto from Dagstuhl Perspectives Workshop 11061

Online Privacy: Towards Informational Self-Determination on the Internet

Edited by

Simone Fischer-Hübner[1], Chris Hoofnagle[2], Ioannis Krontiris[3], Kai Rannenberg[4], and Michael Waidner[5]

1 **Karlstad University, Sweden**
simone.fischer-huebner@kau.se
2 **University of California, Berkeley, U.S.A.**
choofnagle@law.berkeley.edu
3 **Goethe University Frankfurt, Germany**
ioannis.krontiris@m-chair.net
4 **Goethe University Frankfurt, Germany**
kai.rannenberg@m-chair.net
5 **TU Darmstadt, Germany**
michael.waidner@sit.fraunhofer.de

Abstract

While the collection and monetization of user data has become a main source for funding "free" services like search engines, online social networks, news sites and blogs, neither privacy-enhancing technologies nor its regulations have kept up with user needs and privacy preferences. The aim of this Manifesto is to raise awareness for the actual state of the art of online privacy, especially in the international research community and in ongoing efforts to improve the respective legal frameworks, and to provide concrete recommendations to industry, regulators, and research agencies for improving online privacy. In particular we examine how the basic principle of informational self-determination, as promoted by European legal doctrines, could be applied to infrastructures like the internet, Web 2.0 and mobile telecommunication networks.

Seminar 06.–11. February, 2011 – www.dagstuhl.de/11061
1998 ACM Subject Classification K.4.1 Privacy
Keywords and phrases Online Social Networks, Informational Self-Determination, Privacy Enhancing Technologies, Data Protection Directive
Digital Object Identifier 10.4230/DagMan.1.1.1

Online Privacy: Towards Inform. Self-Determ. on the Internet, *Dagstuhl Manifestos*, Vol. 1, Issue 1, pp. 1–20
Editors: S. Fischer-Hübner, C. Hoofnagle, I. Krontiris, K. Rannenberg, and M. Waidner

Dagstuhl Manifestos
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Executive Summary

Existing conceptions of privacy typically incorporate user control as a key component, or indeed describe privacy as a form of user control over information. However, the architecture and development of the Internet have driven individuals to lose control over the collection, use and transfer of their personal data online. Instead, the fundamental value exchange underlying the Internet economy is that services are provided free of charge in return for pervasive use of individuals' information. This business model remains opaque to many users, who willingly or unwillingly share massive amounts of personal data, with a myriad of parties online.

State of the Art. The subjective and contextual nature of privacy challenges any attempt to crystallize individual attitudes toward issues like information sharing on social networks or online behavioural advertising. Perceptions of privacy vary across cultures; there is often inconsistency between what people say about privacy and the options available to them to express their preferences; and privacy harms are difficult to measure. This complicates the understanding of social norms with respect to online privacy.

Privacy enhancing technologies (PETs), while existing for many years, have not been widely adopted by either industry or users. PETs include opacity tools, intended to "hide" personal data in accordance with the principle of data minimization, as well as transparency enhancing tools (TETs), providing users with information about privacy policies or granting them online access to their personal data.

Problems with PETs include low demand from both industry, which is fuelled by information, and users, whose awareness and interest are low; as well as lack of mechanisms for integration into large infrastructures, which were conceived and designed without privacy in mind, requiring properties like usability, scalability, efficiency, portability, robustness, preservation of system security, and more. Problems with TETs include the difficulty and lack of interest of individuals to comprehend complex data flows; and security problems arising from the provision of online access to personal data.

Engineering and Industry Options. Businesses have insufficient incentive to integrate privacy into the design and management of products and services, due to low demand from users; low awareness on the part of both users and industry; competitive business pressures to exploit information; and an absence of coherent, harmonized global regulation.

To improve the current environment, we must meet three challenges: First, transparency must be enhanced for users, through implementation of TETs in a privacy-friendly manner, open source development of TETs, and better user interfaces for transparency in complex environments. Second, PETs should be designed and delivered to end users by building blueprints and sample prototypes for key scenarios (e.g., delivery of service on mobile devices; or use of pseudonyms on communication networks), and deploying open source code to reduce market entry costs. Third, identity management should be promoted as a key technique to manage information while satisfying principles of data minimization and transparency; using minimum data to authenticate and authorize users; and giving particular emphasis to user centric identity management.

Improving Regulations. Unfortunately, neither the current European legal framework nor the United States approach of industry self-regulation has been effective in protecting privacy online. The main problems inflicting the current framework are the blurring distinction between personal and non-personal data; the erosion of consent as a sound basis for data processing; the (in)applicability of European law to websites and third parties based in

the United States; and the regulatory emphasis on ex post remedies in lieu of ex ante risk minimization.

Recognizing that privacy is regarded in Europe not only as an individual right but also as a societal good, which underlies values such as democracy, autonomy and pluralism, we must insist on the continued existence of a strong European legal instrument based on principles of data minimization and ex ante risk prevention. While the path of least resistance may be to make incremental changes to the Data Protection Directive, this may not succeed (and may result in slightly better but still ineffective regulation) absent rectification of fundamental conceptual shortfalls. First, recent examples of de-anonymization attacks have proven the futility of trying to distinguish between personal and non-personal data. Second, given the societal value of privacy as well as the inherently suspect nature of consent in many settings, the limits of consent must be clearly delineated preventing the use of watered down consent to legitimize intrusive processing activities. Third, policy makers should engage with industry not only through lobbyists and trade associations, which pursue a maximalist anti-regulatory agenda, but also with technical experts, system designers, computer scientists and engineers, whose approach towards privacy is more balanced. Fourth, given the global nature of the market for information and ubiquity of cross-border data flows, international enforcement must be coordinated by a central authority, advised by the Article 29 working party; and national privacy regulators should be staffed with not only lawyers but also computer scientists, economists, political scientists, and more, to veer away from their current bureaucratic culture and develop state of the art technological competence.

Additional principles that must be better enforced are privacy by design, requiring comprehensive and iterative privacy impact assessments and implementation of PETs; transparency, providing users with online access to their personal data conveniently, securely, privately, and free of charge, including through the use of "privacy agents;" and accountability, meaning not only passive logging of activity but also the proactive policing and deterrence of abuse within organizations.

Recommendations for Research. First, we suggest research is undertaken to examine the deployment, integration and scaling of PETs in large open-ended networks with decentralized governance and control structures. Research should be multidisciplinary and grounded on empirical data documenting information flows in cloud based applications, ubiquitous computing, and online behavioural targeting. Second, research is needed to promote privacy friendly system engineering, including the transformation of privacy impact assessments from an art into a systematic and transparent process; as well as the integration of PETs through the entire protocol stack via a number of applications, engineering privacy into complete systems and examining methods of evaluation, criteria and metrics. Third, research should seek creative, innovative tools, such as "virtual care-takers," to empower users by enhancing transparency and informational self-determination. Finally, research should explore the "known unknowns", anticipating possible changes to the technological and social environment, such as the impact of quantum computing on cryptographic technologies and the availability of robust face recognition technologies and natural language processing.

Table of Contents

1 Introduction

The principle of informational self-determination is of special importance for online privacy due to the infrastructural and interactive nature of modern online communication and to the options that modern computers offer, even though it is much older than the notion of "Online Privacy". Well before the advent of Web 2.0, the term informational self-determination originated in the context of a German constitutional ruling, related to the 1983 census, making Germany the first country to establish the principle of informational self-determination for its citizens. The German Federal Constitutional Court ruled that[1]: "[...] in the context of modern data processing, the protection of the individual against unlimited collection, storage, use and disclosure of his/her personal data is encompassed by the general personal rights of the [German Constitution]. This basic right warrants in this respect the capacity of the individual to determine in principle the disclosure and use of his/her personal data. Limitations to this informational self-determination are allowed only in case of overriding public interest."

To put it simply, this provision gave individuals the right to determine what personal data is disclosed, to whom, and for what purposes it is used. Informational self-determination also reflects Westin's description of privacy as "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others" [13]. Despite this legal development, the path of the Internet did much to undermine these values and nowadays, individuals have effectively lost control over the collection, disclosure and use of their personal data.

With the evolution and commercialization of the Internet and the advent of Web 2.0, including its search engines and social networks, the environment, in which we need to support online privacy and informational self-determination, became more complex. Some new business models, like ad-financed "free services" for Internet users, rely on a wide-range collection of user data for various purposes, such as marketing of online shops or targeted advertising and include user profiling. It appears that many users of online services are unaware of the implications of this business model. In other contexts, data collected for commercial uses has been later employed for government purposes; this has been possible by the fact that the rules of "free services" place little restriction on reuse of data.

This manifesto is the main result of the Dagstuhl Perspectives Workshop 11061 that took place in February 2011. A primary challenge that it deals with is the correction of power imbalances arising from a loss of informational self-determination, as introduced above. The rest of the document is structured as follows: Section 2 discusses the current state of online privacy, existing technologies to protect privacy, as well as transparency in online information systems, in order for users to have leverage to protect their privacy. Section 3 analyzes what Engineering and Industry can do to improve online privacy. Section 4 gives recommendations on how regulators can improve online privacy, and Section 5 suggests more long-term research topics that are needed to improve online privacy. Finally, the Annex provides further details to these four main chapters.

[1] *BVerfGE 65,1 – Volkszählung*, available in English at http://en.wikipedia.org/wiki/Informational_self-determination.

2 State of the Art

The state of the art in online privacy includes extensive work in a broad spectrum of disciplines. The primary purpose of this section is to set the landscape for the manifesto and provide responses to clusters of specific questions that arise concerning the current state of work in the area of online privacy.

2.1 Understanding Online Privacy

Privacy is subjective, contextual and therefore hard to evaluate. In this regard, one of the main challenges that researchers are currently exploring is linked with the analysis of individual attitudes on privacy. For instance, research has shown that most users of websites with customizable privacy settings, such as Online Social Networks (OSNs), maintain the default permissive settings, which may lead to unwanted privacy outcomes [6]. The explanation to this behaviour is not necessarily that users do not care about their privacy[2]. Instead, existing studies demonstrate an ambivalence of the users' attitudes towards privacy [11, 3]. What makes it more difficult to interpret people's attitude against privacy is that the notion of privacy differs or changes, depending on the culture that individuals are coming from. So, there is still much need for experiments with individuals to allow a broader range of privacy related analysis to be tested and enable a better understanding of people's concerns and the actions they take to address these concerns.

This analysis becomes particularly difficult, since frequently there is no immediate damage for individuals. Even though in some cases, an individual may directly experience an offence, if harassed, manipulated or embarrassed as a result of a prior privacy violation, more frequently the consequences may occur only later or not at all, as for example in third-party tracking of online behaviour for targeted advertisements. Even though tools that try to limit this tracking by third parties exist, with varying degree of effectiveness, these tools do not reveal the impact of processing and usage of personal data from third parties for their own purposes. Evidence also exists to show that third-parties are not only receiving browsing behaviour information about individuals, but are in a position to link this behaviour to personal information via sites such as OSNs [7].

While being subjective and contextual, privacy as a concept has a larger function in society. In this manifesto we discuss privacy under the light of collection and usage practices, but a more general discussion on the basic values that are challenged by the changes brought by a networked society, remains open. What are the effects on our democratic societies of massive-scale data collection, trend prediction and individual targeting? Are people forced into higher conformance? Is conformance pressure affecting the building of political opinions? A scientific approach to these questions cannot rely on the repetition of an old mantra saying that data collection is bad, but will undertake research into the new power relations as they form in the new networked landscape.

[2] On the contrary, several polls and surveys support the opinion that individuals care about their privacy. For example, such a collection for US consumers is presented in http://www.cdt.org/privacy/guide/surveyinfo.php.

2.2 Privacy Technology Landscape and Technology Transfer

In this section, we briefly sketch the current landscape of privacy-enhancing technologies (PETs) and then try to shed light on the reasons for lack of wide-spread adoption of PETs. There is a growing amount of research in the field of PETs, proposing technologies for solving various aspects of the privacy problem; yet PETs are not widely adopted in practice. One cannot expect a simple explanation to this, as online privacy is a complex and interdisciplinary issue. Therefore, we will revisit this issue in the next sections, from the perspective of different disciplines separately with the goal to suggest specific actions. But first, in this section, we set the landscape from a more general view. More specifically, we elaborate on the following reasons, with the understanding that this is an incomplete list of issues:

- current economic environment fosters personal data collection in some business models,
- user awareness of the privacy problems, as well as demand for transparency of data usage and information processing is low,
- today's PETs still lack usability, scalability and portability in many cases,
- regulatory and technical agendas lag behind new data collection practices and data flows,
- integration of many new PETs require costly changes in the existing infrastructure.

After more than 20 years of research in the area of privacy and PETs, there exists a wide variety of mechanisms [4]. Broadly speaking, we could distinguish between opacity tools and tools that enforce other legal privacy principles, such as transparency, security or purpose binding[3]. Opacity tools can be seen as the "classical" PETs, which "hide information", i.e. striving for data minimization and unlinkability. They cover a wide variety of technologies, ranging from cryptographic algorithms and protocols (e.g., [homomorphic] encryption, blind and group signatures, anonymous credentials, oblivious transfer, zero-knowledge proofs etc.) to complex systems like user-centric identity management. Opacity tools can be further characterized depending on whether they focus on data minimization at the network layer or at the application layer. Proposals for achieving sender or recipient anonymity at the network layer comprise protocols such as Chaumian Mixes, DC-Net, etc. At the application layer, a much greater variety of technology proposals exists, such as private information retrieval, privacy preserving data mining (random data perturbation, secure multiparty computation), biometric template protection, location privacy, digital pseudonyms, anonymous digital cash, privacy-preserving value exchange, privacy policies etc.

Transparency-enhancing tools (TETs) belong in the second category of PETs and focus on enforcing transparency, in cases where personal data need to be processed. By transparency we mean the informative representation to the user of the legal and technical aspects of the purpose of data collection, how the personal data flows, where and how long it is stored, what type of controls the user will have after submitting the personal data, who will be able to access the information, etc.

TETs frequently consist of end-user transparency tools and services-side components enabling transparency [10]. The end-user tools include, among other techniques, (1) tools that provide information about the intended collection, storage and/or data processing to the users when personal data are requested from their system (via personalized apps or cookies) and (2) technologies that grant end-users online access to their personal data and/or to information on how their data have been processed and whether this was in line with privacy laws and/or negotiated policies[4].

[3] Purpose binding means that personal data should be relevant to the purposes for which they are to be used and to the extend necessary for those purposes, and should not be usable in other contexts.

[4] A third type of TETs, which has so far only been discussed in the research community, include tools with "counter profiling" capabilities helping a user to "guess" how her data match relevant group profiles,

Examples are the Google Dashboard[5] or the Amazon's Recommendation Service, which grant users online access to their data and allow them to rectify and/or delete their data. However, these are server-side functions and not user-side tools and they usually grant users access only to parts of their data and not to all the data that the respective service processes. An example of user-side transparency enhancing tool is the Data Track developed in the EU project PrimeLife [12], which gives the user an overview of what data have been sent to different data controllers and also makes it possible for a data subject to access her personal data and see information on how her data have been processed and whether this was in line with privacy laws and/or negotiated policies.

In the current state, once the data has been submitted to an online information system, individuals get no knowledge about any further processing. But, even if we assume that the data processing of such complex systems like Facebook, Apple iTunes or Google Search could be transparent to the public, it would be hard or impossible for ordinary individuals to understand what happens with their data. Full transparency of data movements also increases security problems in such environments, if misused with malicious intent. Consequently, this limitation leads to the observation that it is more important for individuals to understand the outcome and implications of data flows in complex online information systems than understanding the full data movements. One technique, among others, that can achieve this kind of transparent outcome-based approach is the creation of ad-preferences by some third-party advertisers, where users are allowed to see the set of outcomes, based on which the data has been forwarded to the third-party (examples here would include Google Ad Categories[6] or the Deutsche Telekom Privacy Gateway for location-based services).

In general, most, if not all, of the proposed PET solutions lag behind the real world situations. They still need to overcome the shortcomings of current approaches, as real world solutions require properties like usability, scalability, efficiency, portability, robustness, preservation of system security, etc. Today, only a patchwork of mechanisms exists, far from a holistic approach to solve the privacy problems. The interaction between these mechanisms and their integration in large scale infrastructures, like the Internet, is not well understood.

Our infrastructures have not been designed with privacy in mind, and they evolve continuously and rapidly integrating new data collection practices and flows. Current privacy mechanisms, not only have difficulties in catching up with these developments, but they also collide with some security and business requirements. A redesign of the system in question can often resolve the collision of interests, but this sometimes requires costly investments.

At the same time, the demand of users for PETs is rather low today. One reason for this is the lack of user awareness with respect to privacy problems, which can be partly attributed to missing transparency of data acquisition and the related information processing, as emphasized above. A second reason lies in the complicated and laborious nature of control imposed on persons, as no legal standards or general consumer protection rules exists. Finally, PETs do not always take into consideration the evolution of privacy models caused by the rapid creation of new technologies and communication models.

Yet another important reason for the lack of adoption of existing PETs lies in some models in data commerce that are based on access to personal data. In the current eco-system, doing nothing about privacy or even aggressively collecting data sometimes pays off, as some companies seem to acquire new clients with new features based on creative data use

which may affect her future opportunities or risks [5].

[5] https://www.google.com/dashboard/

[6] http://www.google.com/ads/preferences

and serendipity. Furthermore, for some players, implementing complex data minimization schemes is costly and time consuming and makes information filtering catered to the end-user much harder, if not impossible. It is important to note here, however, that this approach is not adopted by all industry players. The next section takes a closer look at the problem of adoption of PETs from the industry and suggests addressing specific challenges to overcome this problem.

3 Engineering and Industry Options

Generally speaking, there is a lack of clear incentives for enterprises to manage personal data in a privacy-respecting manner, to design privacy-preserving products, or to make the use of personal data transparent to the data subject[7]. We identify the following root-causes for this current situation:

1. Lack of customer (individuals, business partners) and market demand for privacy respecting ICTs, systems, services and controls (beyond punishments for breaches and other excesses). Usage models for privacy-enhancing technologies cannot currently be targeted to customer demand;
2. Some industry segments' norms, practices and other competitive pressures that favour exploiting personal data in ways contrary to privacy and the spirit of informational self-determination (resulting in diffusion of transparency and accountability);
3. Poor awareness, desire, or authority within some industry segments on the operationalization of privacy (e.g., to integrate existing PETs, to design privacy-respecting technologies and systems, and to establish, measure and evaluate privacy requirements); and
4. Lack of clarity, consistency, and international harmonization in legal requirements governing data privacy within and across jurisdictions (avoided, for example, by migrating data somewhere up in the cloud).

To improve the current environment, we need to increase awareness across users, industry and technologists regarding

- the protection of privacy of users across different media,
- the transparency for processing of personal data,
- the acceptance and incorporation of improved privacy-enhancing technologies by technologists outside of the "privacy community".

To support this goal, we recommend that industry addresses three mid-term challenges, which we discuss in the rest of this section.

3.1 Challenge 1: Promoting Transparency

3.1.1 Transparency-enhancing tools

Transparency enhancing technologies (TETs), which have been developed in the recent years within research projects and by the industry, can help end-users to better understand privacy implications and thus help to increase the user awareness, as we demanded. On the other

[7] We make a disclaimer here that these deficiencies do not apply across the board to all enterprises.

hand, allowing users to control and correct their data processed at services sides will also lead to better data quality for the respective industries.

Challenges for practical TETs that still remain, include the following:

- Providing transparency in a privacy-friendly manner means that TETs should work for pseudonymous users. Industry should consider integration of existing research prototypes and concepts of such privacy-friendly TETs, like the PrimeLife Data Track [8], in real world processes and IT systems.
- The open source development of transparency-enhancing technologies and end-user tools needs to increase, in order to lower market entry costs.
- Better use interfaces for transparency tools in complex environments will need to be created. Also, user-friendly display of data handling practices by "hidden" data processors will play an important role.

3.1.2 Transparency within industrial organizations

Industry needs to foster in-house transparency and awareness for the risks of system-imminent privacy issues in order to effectively enhance privacy in the developed products and services. Principles, such as data minimization and purpose-binding, have to become design principles for processes, IT, service and product design. Industry needs to consistently consider privacy issues, risks, and privacy principles in internal guidelines. These guidelines need to be communicated to engineers, developers, etc. to create a "culture of privacy".

3.2 Challenge 2: Designing and Delivering Privacy Respecting Products to End-users

3.2.1 Demonstrating the power of PETs by blueprints and sample prototypes

When building applications, engineers often lack practical knowledge on incorporating PETs to achieve security and privacy protection. To support engineers in employing privacy-enhancing technologies, we propose to build blueprints and sample prototypes for key scenarios and for different industries. Examples for such prototypes include the following:

- A service that can be delivered to a user on a mobile device, such that the parties involved are able to deliver their parts and are paid for their service, while the user is ensured that every such party receives and stores only minimal data. The user is provided with transparency and control of his own data flows, while data dispersion is minimized, e.g. by attribute-based access-control[8].
- A communication platform that offers its users a convenient communication and collaboration environment with simple and secure user privacy controls to set the audience for certain private data dependent on different social roles and the support of user pseudonyms. The prototype must further demonstrate its economic viability by proper business models that do not conflict privacy requirements.

[8] For example see the ABC4Trust project (https://abc4trust.eu/).

3.2.2 Open or shared-source developments

Sharing source code which can be reused and adopted easily, allows market entrants to lower development costs. One example is the WebKit library[9]. An open-source suite of privacy-enhancing tools can lower market entry costs for companies, which want to offer privacy products and support the emergence of non-commercial software that integrates privacy-protecting functions.

3.3 Challenge 3: Identity Management as a Key Technique

It has been pointed out that identity management is instrumental to the implementation of online privacy management [10]. We also believe that identity management can be used to manage handling of data relevant to satisfy privacy requirements, such as data minimization and transparency.

The scope of identity management is quite broad, comprising authoritative information about legal persons, customer or user relationships, self-issued claims, pseudonyms and anonymous credentials. A minimum of personal data must be conveyed to the service in order to authenticate and authorize the accessing subject.

The service-side storage of personal information without transparent and traceable relation to identities creates fundamental asymmetries in the relationship between the users and the industry and erodes transparency, confidence and trust. Therefore, we propose user-centric identity management systems, which can restore this balance and confidence.

User-centric identity management in this context implies that personal data – even in cases that is created by a service – is always handed back to the user upon completion of the service. If the user desires consistency across service invocation, it is her decision to hand over the data again to the same or another service. This way, individuals can supervise and limit personal data disclosure and exercise rights of access to their data held by third parties.

User-centric identity management allows users to detect any linkages to third parties created from the primary relationship. Enterprise policies and procedures should support user-centric identity management as well, to prevent unwanted linkages and inadvertent disclosures of personal data.

4 Recommendations for Improving Regulations

4.1 Current Regulatory Framework Insufficient

Neither the current European legal framework, nor the US approach toward private sector self-regulation, has been effective for the protection of privacy online, particularly with regard to new business models, such as behavioural targeting, user profiling, social networking and location-based services. Key weaknesses in the EU framework include that: 1) services based predominantly in the US are effectively outside European jurisdiction 2) European users have little choice but to "consent" to companies' terms of use and privacy policies in the absence of alternatives of comparable functionality, 3) the concept of "personal data" is currently the necessary trigger for the applicability of the Data Protection Directive (DPD) and 4) there

[9] http://www.webkit.org/

seems to be too much reliance on ex post securing of data rather than on ex ante elimination of privacy risks through data minimization (for example the recent Art.29 WP Opinion on smart metering[10] omitted entirely any consideration of radical data minimization through cryptographic methods[11]).

4.2 Distinctive European Privacy Values

The European culture of privacy incorporates values of democracy, autonomy and pluralism. The European views on privacy as a societal good and as a factor of public interest lead to a more prominent role of the State in this domain. This conception is not widely shared outside Europe, where the notion of privacy is strongly linked to "the right to be let alone". Consequently, countries such as the US do not necessarily establish the same balance between economical needs and privacy protection.

European approaches protect privacy through consumer protection interventions, instead of reliance upon contract. For instance, it is conceivable that a European national government might prohibit certain extremely privacy-invasive practices, like long-term storage of online search requests for commercial purposes. Unlike contract approaches, such prohibitions can never be waived by acquiring the consent of the users[12]. It has to be recognized that this European view is not shared by legislators in other parts of the world.

4.3 Surveillance Society and Blanket Retention of Data

An important issue of principle for the future Internet of things is the legitimacy of the blanket retention of traffic data (or metadata). In so far as such data relates to individuals, it constitutes a "map of private life" [2]. Case law of the European Court of Human Rights (ECtHR) establishes that "merely" storing such data engages the right to privacy. The troubling exception to this rule is the Data Retention Directive (DRD), requiring storage of certain telecommunications and Internet traffic data. However, the legitimacy of the DRD remains controversial and the concept of indiscriminate continuous retention of data about the entire population has been ruled unconstitutional in its entirety by the Romanian Supreme Court[13], because it "makes the essence of the right disappear".

4.4 A Strong European Legal Instrument Remains Useful

Notwithstanding the fact that a global harmonization in this area is not yet possible, a strong and effective European legal instrument has the potential of having an impact on the global online context. The essential question in conceiving a unique, strong and effective European legal instrument is the goal we want to achieve. The first fundamental objective should be the prevention of privacy-endangering information-processing practices at all levels.

[10] http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2011/wp183_en.pdf
[11] http://research.microsoft.com/en-us/projects/privacy_in_metering/
[12] Art. 8.2, a) of the Directive provides that in certain cases the prohibition to process sensitive personal data may not be lifted by the data subject's giving his consent
[13] http://www.legi-internet.ro/english/jurisprudenta-it-romania/decizii-it/romanian-constitutional-court-decision-regarding-data-retention.html

This instrument must guarantee real protection against actual and potential risks taking into account technological developments and not merely offer formalistic legal assurances. Therefore, it is crucial to take maximum advantage of the opportunity offered by the current proposed review of the European directive to maximize its impact.

The path of least resistance is to make incremental changes in the existing directive. However this may not succeed in rectifying some serious conceptual defects. There is a risk that the review of the directive will merely result in a slightly better but still largely ineffective regulatory solution.

Moreover recent results in the field of de-anonymization [1, 8] suggest that some data may be impossible to anonymize (e.g. social networks) and it is difficult to predict the vulnerabilities for and consequences of re-identification when contemplating the release of pseudonymous data [9]. It seems that a better approach would be to make the application of the legal framework dependent on an evaluation of the actual and potential privacy risks related to any data processing.

4.5 Consent Must not Overrule Everything

From a European perspective, explicitly given user consent should not be accepted as a waiver for privacy-intrusive online practices. A European legal instrument should clearly emphasize that individual privacy is not purely a matter of the individual concerned, but of the society as a whole. Moreover in many situations the voluntariness of the user's consent can be put into question because of the lack of reasonable alternatives for commonplace services, which meaningfully adhere to European Data Protection. US consumer protection law recognizes many situations where consumer consent cannot waive risks. This policy has not yet mobilized into the law of privacy. Consent should expire, according to the scope and extent of processing. When asking for consent, data controllers should make explicit what is revocable and what is irrevocable and how it is possible to revoke that consent. Legislation may prohibit processing that would have irrevocable consequences. Otherwise a European-wide warning system for specific risks or breaches, similar to governmental travel warnings for dangerous regions, could be advisable.

4.6 Effective Implementation and Enforcement Is Crucial

Privacy lobbying has been concentrated amongst a few law firms and trade associations. These interests pursue a maximally anti-regulatory agenda, even in situations where their clients would admit that they could comply with privacy rules proposed. Retention of information in network advertising companies is a key example – while on the lobbying level, it is often argued that this data must be retained for very long periods of time, the engineers at network advertising companies will admit that data becomes less valuable after a very short period of time, and is often unused for ad targeting purposes within a few months. However, since policy makers rarely engage beyond trade associations, they get a jaundiced view of the actual requirements that businesses have for data. Too often, regulators pose technical and implementation questions to attorneys rather than the technical experts who design systems. We recommend that to the extent practicable, regulators invite relevant actors, rather than their representatives, to public fora, consultations, and other consensus-building events around privacy.

The US Federal Trade Commission recently employed two technical experts on a short term basis to assist in the evaluation of technologies, and the agency has, also on a short term basis, employed a senior computer scientist to assist with policy analysis. We believe that technical expertise is increasingly necessary for policy makers and regulators, and we recommend they look more to in-house technical expertise to assist in their rule-making and investigations.

So far the DPD agenda for reform has not sufficiently considered basic limitations on the effectiveness of enforcement when 27 national authorities must reach consensus. Competence for international enforcement actions should be given to a central authority, advised by the Art. 29 WP, leaving national DPAs better able to focus on national-level issues. Current technological knowledge must be an indispensable part of the professional competence of DPA administrations. But at most a few percent of these officials have any relevant postgraduate scientific competence, and overwhelmingly DPAs have an irredeemably bureaucratic culture. A complete renewal of these institutions is necessary. A minimum of one third of staff should be experts in the computer science of privacy, as well as first rate talent in law, economics, political sciences, sociology and philosophy. Access to justice through privacy litigation is out of reach for most people today. Data Protection Authorities should evolve into Information Privacy Ombudsman (IPO), explicitly acting to uphold privacy rights. IPOs must expect to show intellectual excellence in every relevant field, and earn their authority through merit, or be dissolved.

4.7 Privacy by Design

A privacy by design approach can be mandated (or otherwise encouraged) by legal or regulatory provisions, if scientific discoveries demonstrate that a service can be offered practicably in a more privacy protecting way. This could involve, for example, requiring that comprehensive and iterative privacy risk and impact assessments be carried out and that state-of-the-art privacy technologies be adopted.

4.8 Transparency for Data Subjects

In order to give meaningful effect to the right to informational self-determination, it is clearly necessary for users to have the possibility of "information self-awareness". Its importance has been emphasized in all previous sections already, together with the limitation in the corresponding transparency enhancing tools. Because invoking existing "subject access" rights is cumbersome, slow, and often incomplete, these rights should be strengthened to provide a right to comprehensive online access to data about an individual's use of an online service, conveniently, securely, privately, and free of charge. This should henceforth be regarded as an indispensable aspect of the human right to privacy in the 21st century. To provide such data genuinely in "intelligible form", more disclosure of algorithms will be necessary (also for automated processing or anonymization), whether these act on personal data or can affect the individual through the application of statistical data models ("red-lining").

Consumers' ability to designate "privacy agents" as proxies for exercise of their rights should be recognized by firms and governments. Consumer privacy agents are now a viable business, but they are frustrated by organizations that question the authority of the agent to act for the consumer, and by systems that attempt to obfuscate the invocation of rights

to opt out or gain access to personal information. Collective negotiation through "privacy unions" potentially is also an important democratic mode of political expression and must be protected from harassing lawsuits.

4.9 Transparency by Design for Auditors

A further important aspect of transparency is the need to design mechanisms, which allow the flows of data in a system to be documented and verified by internal and external auditors, including algorithms used to perform profiling and social sorting.

4.10 Accountability

A core reason for defining the notion of a data controller was to assign clear legal responsibility for compliance. However the complex mesh of legal relationships, which have since arisen, often do not allow a controller to guarantee any effective operational performance of such obligations. Mere logging of system activity is insufficient to counter insider threats – active policing of such logs is required. "The Principle of Accountability" should be understood to mean not merely a passive ability to provide an account, but the creation of an effective deterrent against abuse. Moreover the creation of detailed logs about data subjects itself is prejudicial to privacy and therefore all logging activity must be assessed from the point of view of the interests of privacy protection as well as justifiable security goals.

5 Recommendations for Research

The up-scaling of privacy-enhancing technology to larger systems and its integration with existing systems fails, mainly because systems aspects and the related interdisciplinary issues are not taken into account. In this section we address this by recommending research into:

- scalability and integration on a large scale,
- technologies to support privacy-enhanced systems engineering and
- research to enable systems for individual-level privacy protection.

Finally, we recommend research into the "known unknowns" of the technological and societal environment that privacy technology exists in.

5.1 Web-Scale Integration, Deployment and Infrastructures

As discussed in Section 2, over the last 20 years, the privacy community has developed a large pool of tools and primitives. Yet, we do not see deployment in large scale infrastructures such as the Internet and the World Wide Web, and the interaction between individual technological tools is ill-understood.

To address this limitation, we recommend research and experimentation focusing on integration and deployment: How do privacy-enhancing technologies scale, in terms of deployment on large and open-ended networks with decentralized control and governance structures, large populations, and qualitatively different scales of data collection and processing? Research instruments to address this question may include: mathematical modelling of interdependencies

and integration effects, test beds and demonstrators that enable research and demonstration, as well as private-public partnerships focusing on adoption and deployment of experimental technologies. This approach can also foster infrastructure and product development through pre-commercial procurement. Also other incentives for deployment should be analysed. Specific fields, in which these approaches should be tried, include (but are not limited to):

- Privacy-enhanced identity management infrastructures;
- Techniques for minimal data disclosure;
- Data governance and policy language approaches;
- Accountability in data disclosure and processing, including transparency and auditability, as well as real-time detection and investigation of privacy breaches and data abuse;
- Technological approaches that help to reconcile privacy interests and business models;
- Privacy-protected communications.

Research approaches towards these questions need to be multidisciplinary. Relevant disciplines include economics, psychology, sociology, business administration, law and political studies, as well as various fields within the discipline of computer science (ref. Annex for examples).

Within this research agenda of understanding large-scale, system-level interactions of technological and social phenomena, the empirical data about the evolution of data collection practices and data flows on the Internet and the Web become a critical asset. Relevant data flows include data treated by cloud-based applications, sensors that interact with the physical environment and users' behaviour as they interact with online services. Regulatory and technical agendas need to be informed by empirical understanding of these data flows. We recommend creating an observatory for these flows and interdependencies, taking existing research work to a systematic new level, and creating the basis for more rigorous analysis of the technical *status quo*.

5.2 Towards Privacy-friendly System Engineering

The privacy enhancing building blocks available today need to be integrated into an overall privacy engineering environment, so as to enable adequate evolution of privacy concepts and the required privacy friendly technology and systems in the future. Requirements for privacy need to be analysed, especially when new systems are coming up that can have a negative impact on privacy. Consequently, a Privacy Impact Assessment (PIA) is needed. PIA requires research about methods to develop it from an art into a systematic and transparent process, which also allows the comparison of different development alternatives and their privacy impact. When privacy enhancing technologies are deployed on servers, network infrastructures and devices, multiple independently developed technologies are brought together. The way in which the privacy properties of these modules interact with each other and with the surrounding system through the entire protocol stack and via a number of applications, is often ill-understood. For example, integration of different systems can lead to surprising effects (e.g., unwanted data flows) resulting in the violation of privacy policies or assumptions implicit in privacy technologies.

Further research in the composability (e.g. considering the current research in "differential privacy") of these tools and systems is needed to develop suitable best practices. First, this research will contribute to the development of methodologies and guidelines that contribute to the ability to engineer practical privacy in complete systems with the help of a multi-stakeholder community. In order to evaluate the privacy assurances given by these approaches,

further research into evaluation methods, criteria and metrics is necessary. Second, the research direction proposed here will also facilitate re-engineering processes of deployed systems to take privacy aspects into account.

5.3 Individual Protection

In the area of individual protection, we can frame many privacy concerns in terms of power imbalances between data subjects and data processors, and we can frame privacy enhancing technologies as tools to assure or restore an adequate power balance. Research should continue towards tools that assist individuals' informational self-determination and permit users to learn, e.g. when they share data and may not know about the consequences. Those tools should leverage progress in machine learning. We could imagine relevant tools ("care-takers"), as for example:

- advisers, helping users before they engage in privacy-relevant activities online,
- bodyguards, assisting users as they act online,
- litigators that might be able to help users reconcile breaches of their expectations afterwards.

A crucial element of individual protection and autonomy is individuals' ability to understand and act on their context, assisted by appropriate and intuitive tools. Related to the observations on scaling in the previous sections, research should address how the implications of massive-scale data collection and processing can be made comprehensible and practically manageable for individuals. Research topics here range from usability of technology to the development of philosophical and psychological models for the consequences of data processing. Additionally, empirical experiments should be designed to better understand what users' privacy interests and assumptions are, and to what extent they are (or are not) able to take action using the tools available today.

5.4 Known Unknowns: Possible Changes to the Technological and Societal Environment

Research agendas in privacy need to address the evolution of underlying technologies and the surrounding societal and business landscape, in particular in cases where that evolution might create qualitative changes to the privacy landscape. Cryptographic technologies build the foundation for controlling access to data and are also used as a primitive in many privacy enhancing tools. When there are risks that cannot be articulated, such as whether Quantum Computing will lead to negative consequences for cryptographic primitives as available today, the question is raised whether we are prepared for the consequences of changing the underlying assumptions that today's technology is built on.

Other examples can be found in the rapid advance in the availability and quality of face recognition technology and natural language processing. Some of these progresses are further aided by the increasing availability of large data sets. The implications of these effects are likely compounded by the availability of more powerful mobile devices. Finally, we should also include unpredicted events and disasters to the factors that may change societal attitudes toward privacy in the future. More generally, blue-sky research should be undertaken to identify and prepare for changes in underlying technologies and broader science and societal context that we might not foresee today.

Annex A: Examples for research approaches on new privacy technologies

As mentioned in Section 5, research approaches on new privacy technologies should be multidisciplinary. Examples for the principles and use cases to check for multidisciplinary questions include:

- Cryptographic feasibility
- Scalability
- Usability/acceptability
- Regulatory
- Business models. Is the new technology approach compatible with the future?

For example, the above principles should be checked on the following upcoming technologies in privacy preserving or privacy-friendly distributed systems and activities:

- Privacy preserving distributed data mining and processing. In this category falls for example the application of Peer-to-Peer architectures on online social networks, in order to avoid control over user data and behaviour by a single entity, such as the service provider. Another example could be technologies targeting the protection against Spam or DDoS, in existing anonymous transport networks.
- Privacy preserving distributed data collection. In this category falls for example the sensing and collection of environmental data that are connected with the context of specific people (e.g., their location). This is the case, when sensors embedded in mobile devices are used for such a collection. While this is an upcoming technology, the privacy implications have been hardly studied.

Annex B: Participants and Observers

Participants

- Andreas Albers
Goethe University
Frankfurt, DE
- Caspar Bowden
Microsoft WW Technology
Office, GB
- Sonja Buchegger
KTH Stockholm, SE
- Johannes A. Buchmann
TU Darmstadt, DE
- Jacques Bus
Digitrust.EU – Brussels, BE
- Jan Camenisch
IBM Research – Zürich, CH
- Fred Carter
IPC – Toronto, CA
- Ingo Dahm
Deutsche Telekom AG, DE
- Claudia Diaz
K.U. Leuven, BE
- Jos Dumortier
K.U. Leuven, BE
- Simone Fischer-Hübner
Karlstad University, SE
- Dieter Gollmann
TU Hamburg-Harburg, DE
- Marit Hansen
ULD SH – Kiel, DE
- Jörg Heuer
Deutsche Telekom AG
Laboratories, DE
- Stefan Köpsell
TU Dresden, DE
- Ioannis Krontiris
Goethe University Frankfurt, DE
- Michael Marhöfer
Nokia Siemens Networks –
München, DE
- Andreas Poller
Fraunhofer SIT – Darmstadt, DE
- Kai Rannenberg
Goethe University Frankfurt, DE
- Thomas L. Roessler
W3C, FR
- Kazue Sako
NEC, JP
- Omer Tene
Israeli College of Management
School of Law, IL
- Hannes Tschofenig
Nokia Siemens Networks –
Espoo, FI
- Claire Vishik
Intel – London, GB
- Michael Waidner
TU Darmstadt, DE
- Rigo Wenning
W3C / ERCIM, FR
- Alma Whitten
Google London, GB
- Craig E. Wills
Worcester Polytechnic Inst., US
- Sven Wohlgemuth
National Institute of Informatics –
Tokyo, JP

Observer

- Jesus Villasante
European Commission, BE

References

1 Lars Backstrom, Cynthia Dwork, and Jon Kleinberg. Wherefore art thou r3579x?: anonymized social networks, hidden patterns, and structural steganography. In *Proceedings of the 16th International Conference on World Wide Web (WWW '07)*, pages 181–190, Banff, Alberta, Canada, 2007.

2 C. Bowden. Closed circuit television for inside your head: Blanket traffic data retention and the emergency anti-terrorism legislation. *Computer and Telecommunications Law Review*, March 2002.

3 L. Brandimarte, A. Acquisti, and G. Loewenstein. Privacy concerns and information disclosure: An illusion of control hypothesis. In *Proceeding of the 9th Workshop on the Economics of Information Security (WEIS 2010)*, June 2010.

4 G. Danezis and S. Gürses. A critical review of 10 years of privacy technology. In *Proceedings of Surveillance Cultures: A Global Surveillance Society?*, London, UK, April 2010.

5 Mireille Hildebrandt. Behavioural Biometric Profiling and Transparency Enhancing Tools. Technical Report FIDIS Deliverable D7.12, March 2009.

6 Balachander Krishnamurthy and Craig E. Wills. Characterizing privacy in online social networks. In *Proceedings of the 1st Workshop on Online Social Networks (WOSN '08)*, pages 37–42, 2008.

7 Balachander Krishnamurthy and Craig E. Wills. On the leakage of personally identifiable information via online social networks. *ACM SIGCOMM Computer Communication Review*, 40:112–117, January 2010.

8 Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *Proceeding of the IEEE Symposium on Security and Privacy*, pages 111–125, 2008.

9 P. Ohm. Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, 57:1701, 2010.

10 K. Rannenberg, D. Royer, and A. Deuker, editors. *The Future of Identity in the Information Society – Challenges and Opportunities*, page 507. Springer, 2009.

11 J. Turow, C. J. Hoofnagle, D. K. Mulligan, N. Good, and J. Grossklags. The federal trade commission and consumer privacy in the coming decade. *I/S: A Journal of Law & Policy for the Information Society*, (723), 2007–08.

12 J. E. Wästlund and S. Fischer-Hübner. End User Transparency Tools: UI Prototypes. Technical Report PrimeLife Deliverable D4.2.2, June 2010.

13 A. Westin. *Privacy and freedom*. New York: Atheneum, 1970.

Manifesto from Dagstuhl Perspectives Workshop 10482

Modeling, Analysis, and Verification – The Formal Methods Manifesto 2010

Edited by

Jörg Kreiker[1], Andrzej Tarlecki[2], Moshe Y. Vardi[3], and Reinhard Wilhelm[4]

1 **TU München, Germany,** `kreiker@in.tum.de`
2 **University of Warsaw, Poland,** `tarlecki@mimuw.edu.pl`
3 **Rice University, U.S.,** `vardi@cs.rice.edu`
4 **Saarland University, Germany,** `wilhelm@cs.uni-saarland.de`

Abstract

This manifesto represents the results of the Dagstuhl Perspectives Workshop 10482 *"Formal Methods – Just a Euro-Science?"* held from November 30 to December 3, 2010 at Schloss Dagstuhl, Germany. We strive to clarify the terminology and categorize the abundance of concepts and methods in order to reduce misunderstandings prevalent in the research community and in communication with industry. We discuss the industrial acceptance of formal methods and how to increase it by targeted research and improved education. Finally, we state a few challenges and provide perspectives of the field.

This document is opinionated in nature and biased towards the experiences and views of the participants listed in the appendix, further distilled by the authors.

Seminar 30. November–03. December, 2010 – www.dagstuhl.de/10482
1998 ACM Subject Classification B.4.4 Performance Analysis and Design Aids, D.2.4 Software/Program Verification, F.3.1 Specifying and Verifying and Reasoning about Programs
Keywords and phrases Formal methods, Verification, Analysis, Modeling, Design for Verifiability
Digital Object Identifier 10.4230/DagMan.1.1.21

Formal Methods – Just a Euro-Science? (10482), *Dagstuhl Manifestos*, Vol. 1, Issue 1, pp. 21–40
Editors: Jörg Kreiker, Andrzej Tarlecki, Moshe Y. Vardi, Reinhard Wilhelm

Dagstuhl Manifestos
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Executive Summary

Formal methods are employed during system-development process to improve the quality of the system, to increase the efficiency of the development process, or to derive guarantees about qualities of the system. The term *"formal methods"* has traditionally been used for a number of different approaches, including *modelling* and *specification languages*, as well as *methods and tools to derive properties of systems*. Because of the vagueness of the term "formal methods", it may perhaps, be desirable to replace it by *"modelling, analysis, and verification"*.

A good recent overview of industrial projects concentrating on the early phases of specification and design has been given in a recent survey article: Jim Woodcock, Peter Gorm Larsen, Juan Bicarregui, John S. Fitzgerald: Formal methods: Practice and experience. *ACM Computing Surveys*, 41(4), 2009.

The Dagstuhl Perspectives Workshop, held in December 2010, concentrated mostly on methods for system analysis and verification. These are employed in the design phase as well as in later phases of system development. Model checking, abstract interpretation, equivalence checking, and verification by deduction –all developed in academia– are the most impressive success stories.

After a very long gestation period, formal methods for the derivation of program properties have finally gained some measure of industrial acceptance. There are, however, remarkable differences in the degree of this acceptance. There is a clear correlation between the criticality of systems and the costs of failure, on one hand, and the degree to which formal methods are employed in their development, on the other hand. Hardware manufacturers and producers of safety-critical embedded systems in the transportation industry are examples of areas where applications of analysis and verification methods are perhaps most visible. A semiconductor design gone wrong is just too costly for any cost argument against the use of formal design and verification tools be acceptable. Threats of liability costs are strong arguments for the use of formal methods in the development of safety-critical embedded systems. Different application areas often entail different approaches to the use of formal methods. Safety-critical systems call for the use of sound methods to dogmatically ensure correctness. General-purpose software with strong time-to-market pressures encourage a more pragmatic attitude, with emphasis on bug-chasing methods and tools.

Industrial domains with certification requirements have introduced tools based on formal methods into their development processes. Most current certification regulations are, however, still process-based; they regulate the development process and do not state the required properties of the result. Critics describe this as "Clean pipes, dirty water." The trend to use formal methods will become stronger when certification standards move from process-based assurance to product-based assurance. These new standards will specify the guarantees to be given about system properties. Several current standards for transportation systems highly recommend abstract interpretation and model checking for systems at the highest criticality level. "Highly recommend" actually means "required". The loophole is the "state-of-practice" argument. The developer can be exempted from using a highly recommended method by arguing that it is not yet the state of practice.

Several participants of the workshop have expressed the important role of champions of a formal method. A champion, enthusiastic about the potential of and competent in the use of a verification method, is often needed to introduce the method and associated tool to the development process. Often, once the champion leaves, the degree of adoption declines dramatically.

The expectations towards analysis and verification methods have always been very high, often due to unrealistic promises. These unrealistic promises have mostly been the result of the ignorance of the differentiation of roles. Three distinct *roles* are connected to a formal method: the *researcher* develops the theoretical foundations of the method; the *tool developer* implements the method; and the *users* apply the tool in an industrial setting. The different analysis and verification methods have very different requirements imposed on their users, which has implications for their acceptance in industry. Researchers and tool developers often develop their methods and tools for their own use. Subsequently, they use these tools with a high degree of expertise. The experience of such expert users is quite different from that of industrial users, who do not have such degree of expertise. Thus, reports by expert users are often quite rosy and create unrealistic expectations. The expectations towards analysis and verification methods are astonishing in the light of the known undecidability or intractability of the problems they are expected to solve; the methods and tools are expected to be at the same time fully automatic, effective and efficient, and easy to use. Disappointment is unavoidable. Nevertheless, the border between what can currently be done and what is still out of reach is permanently moving, with significant progress accomplished over the last 30 years.

One challenge for further advances is higher *degree of automation*: the different methods require different degrees of user interaction and of user qualification. Currently, with few exceptions, such as Microsoft Research's Boogie platform, there is little integration among different tools. Nevertheless, advances can be expected in the coming years from *tool integration*, starting with information exchange between tools and common exchange formats. Specifically, there is a high potential for improvement from a synergetic integration of model-based design tools with analysis and verification tools.

Scalability of the methods and tools is still considered a problem. The *exploitation of large-scale parallelism* may increase the size of verifiable systems. A clear *identification of application areas* for the various methods rather than the search for universal methods, doomed to fail, will avoid user disappointment.

The embedded-system industry has already realised that badly structured systems written in obscure programming style cannot be effectively maintained. Similarly, it cannot be expected that verification methods would cope with such systems. Systems should be *designed for verifiability*.

While formal methods have often been dismissed by many as "Euro-Science" –a rather abstract research with little chance for industrial adoption– decades of research, both basic research and tool development have started to bear fruits, attracting an increasing level of industrial interest. This interest is often accompanied by unrealistic expectations, but, at the same time, provides an opportunity and challenge to researchers working in this area, as more basic research and good tools engineering are needed to solve the challenges outlined above.

Table of Contents

1 Introduction

Formal methods are employed during system-development process to improve the quality of the system, to increase the efficiency of the development process, or to derive guarantees about qualities of the system. After a very long gestation period, formal methods for the derivation of program properties have finally gained some measure of industrial acceptance. There are, however, remarkable differences in the degree of this acceptance. There is a clear correlation between the criticality of systems and the costs of failure, on one hand, and the degree to which formal methods are employed in their development, on the other hand. Hardware manufacturers and producers of safety-critical embedded systems in the transportation industry are examples of areas where applications of analysis and verification methods are perhaps most visible. Different application areas often entail different approaches to the use of formal methods. Safety-critical systems call for the use of sound methods to dogmatically ensure correctness. General-purpose software with strong time-to-market pressures encourage a more pragmatic attitude, with emphasis on bug-chasing methods and tools.

The expectations towards analysis and verification methods are astonishing in the light of the known undecidability or intractability of the problems they are expected to solve; the methods and tools are expected to be at the same time fully automatic, effective and efficient, and easy to use. Disappointment is unavoidable. Nevertheless, the border between what can currently be done and what is still out of reach is permanently moving, with significant progress accomplished over the last 30 years.

One challenge for further advances is higher *degree of automation*: the different methods require different degrees of user interaction and of user qualification. There is little integration among different tools. Nevertheless, advances can be expected in the coming years from *tool integration*, starting with information exchange between tools and common exchange formats. Specifically, there is a high potential for improvement from a synergetic integration of model-based design tools with analysis and verification tools.

Scalability of the methods and tools is still considered a problem. The *exploitation of large-scale parallelism* may increase the size of verifiable systems. A clear *identification of application areas* for the various methods rather than the search for universal methods, doomed to fail, will avoid user disappointment.

The embedded-system industry has already realised that badly structured systems written in obscure programming style cannot be effectively maintained. Similarly, it cannot be expected that verification methods would cope with such systems. Systems should be *designed for verifiability*.

While formal methods have often been dismissed by many as "Euro-Science" –a rather abstract research with little chance for industrial adoption– decades of research, both basic research and tool development have started to bear fruits, attracting an increasing level of industrial interest. This provides an opportunity and challenge to researchers working in this area, as more basic research and good tools engineering are needed to solve the challenges outlined above.

2 Concepts

Defining formal methods is easy. I did it 100 times.[1]

[1] Adapted from a quote by Mark Twain on quitting smoking.

We understand the area of *modelling, analysis and verification* to be mathematically founded *techniques* and *tools* that aid *humans* to *construct systems* of a higher *quality* with less *resource* usage.

Systems

Systems we consider consist of software, hardware, or a combination thereof. We use the term *data* to abstract the interaction of systems with the physical world (such as input from users or sensors). A system together with data *executes*, resulting in a (number of) mathematically defined *run(s)* of the system.

Verification

(Functional) Verification is concerned with the behavior *intended by the constructor*: 'A system should do, what I want it to do' Intended behavior is a set of acceptable runs. Acceptable runs can be defined implicitly or explicitly. Explicit definitions of acceptable runs are called *specifications*. Typically, specifications are (parts of) programs written in a formal language. The *verification* of a system formally proves its conformance with its *specification*.

Verification is the most abused term in this arena. Almost any activity to convince oneself or a client of desired system properties is called *verification* or even *formal verification*. Most notably (non-exhaustive) testing and bug finding sail under this false color. What contributes to this confusion is the polyvalent nature of several formal-method approaches. Model checking, advertised as a verification technique, is, for complexity reasons, mostly used for bug finding. We will use *verification* in the strict sense, defined in Subsection 3.2.6.

Resources

We distinguish two kinds of resource. On the one hand, there are resources consumed during the *construction* and for the *maintenance* of a system. On the other hand, there are resources consumed during the *runs* of the system. Examples of resources are time, energy, person months, and money.

Roles

Humans are involved with formal methods in several roles.

- *researchers* develop the foundations of a formal method,
- *tool developers* realize a tool based on some formal foundation,
- *users* apply tools to systems under study.

Some tools are push-button tools. However, given the undecidability of the program verification problem, more or less input from the tool user is needed as will be detailed when individual methods are described. The formal-methods area has long suffered from the fact that the researcher was also the tool developer, and worse, the user. This type of user has made unusually good experience with his own method and tool, and often raised too high expectations.

Models

Systems are the end-product of a construction process. A *model* is then viewed as an abstraction of a system. It should be formally defined (possibly in the same language as the

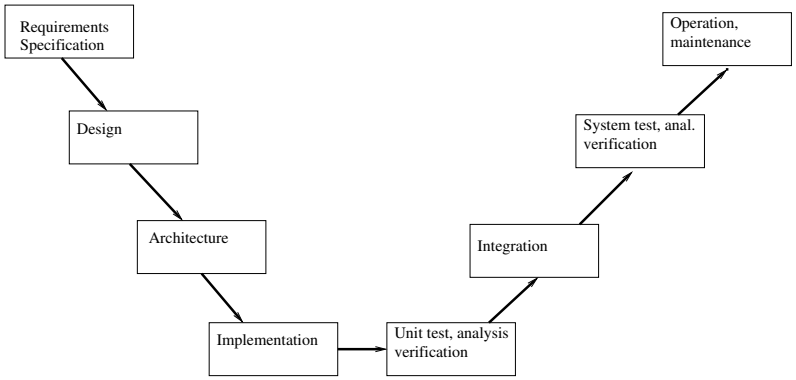

Figure 1 The V cycle for software development

system). Formal methods study the relation between a model and the system *implementing* the model.

Having introduced general definitions, we instantiate them with existing systems, specifications and methods in the next section.

2.1 Development Process

The construction of systems is organized according to some development process. One such process is depicted in Figure 1. Typically, different formal methods contribute to that process and, typically, each step during the development is associated with preferred methods.

Modeling and specification languages are associated with the first three phases, Requirements Specification, Design, and Architecture. Code synthesis in model-based approaches is associated with Implementation. Model checking, abstract interpretation, and program verification may be appplied at the model level during the Design phase and at the unit and the system level.

3 A Survey of Formal Methods

In this section, we list typical formal methods in no particular order. It must be noted that different application domains require different methods and that there is not one ideal method. Therefore, we delineate typical application domains with each method.

Method Signatures

Following Dines Bjørners proposal at the workshop, we ask developers and users of formal methods to provide a *signature* of their method. A signature is a type expression over the type constructors: Sys, Model, Data, and Spec, $\rightarrow$, and $\times$. As an example consider

$$\mathit{compiler} :: \mathsf{Model} \rightarrow \mathsf{Sys}.$$

That is, a compiler takes a model (here: a program) and produces a system (here: binary code). Note that binary code could serve as a model (for instance in worst-case execution time analysis) and that a program could serve as a system (if one is interested in

the development resulting in the program, say, from a higher specification by model-based development).

3.1 Specification Languages

Specification languages, as the term indicates, are formalisms to write *specifications* that capture properties required of the system (or better, of its behaviour). Therefore, the basic functionality they offer is to built specifications that capture designer's initial, informal ideas about desired system behaviour:

$$\mathit{specify} :: \mathsf{Informal\_Ideas} \to \mathsf{Spec}.$$

Once specifications are given, to make them of any use, it must be possible to determine whether or not a given system *satisfies* a specification, or in other words, *realizes* it correctly:

$$\mathit{correct} :: \mathsf{Sys} \times \mathsf{Spec} \to \mathbb{B}.$$

Checking this relationship is the matter of *verification* of a system w.r.t. its specification.

Refinement is a relation between specifications, whereby a specification gets refined by incorporating some further design and programming decisions, with the basic requirement that any system that correctly realizes the refined specification satisfies the original specification as well:

$$\mathit{refines} :: \mathsf{Spec} \times \mathsf{Spec} \to \mathbb{B}.$$

The approaches to capturing the required system behaviour (or some of its aspects) vary vastly.

At one extreme, we have the use of standard (high-level) programming languages to define a specific behaviour, which is then viewed as a presentation of the desired behaviour of the system. This idea led to the development of *model-oriented* specification languages, where a particular model of a system is built (whether using some programming formalisms, or some more abstract mathematical notation) and then viewed as a specification: any system that displays a behaviour conformant with that of the given model is its correct realization. Development then takes a form of *model refinement*, or reduces to *synthesizing* a program from a model. An archetypical example of a model-oriented specification formalism, which heavily influenced further developments is VDM [20].

Model-based design is extensively used in the embedded-systems industry. It offers a unique chance, namely that code generators are tailored towards analyzability and verifyability. It is not the programmer who has to be forced to produce disciplined code. Code synthesis, as experienced with several code generators for Scade, can be designed to guarantee analyzability and verifiability.

Another corner is occupied by logical formalisms, where a specification is given by simply listing the required properties as formulae of some logic. Specifications for programs-in-the-small are often given in this form, based on a programming logic (with Hoare's logic [19] as a classical example). Not much more than pure formulae of some temporal logics have traditionally been used in model checking. Finally, various algebraic languages were put forward. The purpose was to specify abstract data types by listing axioms that link the operations (and constants) of a data type with each other. They thus specify the behaviour of functions that are to implement the operations [16]. These ideas spurred development of *property-oriented* specification languages, like Z [33], which apart from such basic specifications, built by listing the axioms, offer various mechanisms to combine specification

systematically, thus building them in a structured manner. This was perhaps first stressed in the work on CLEAR [10] and developed in full in a whole line of algebraic specification languages, with CASL [6, 26] as a relatively recent fully-fledged example of such a formalism, with complete formal semantics, development and verification methodologies, and support tools under intensive development.

The distinction between model- and property-oriented approaches to specifications was blurred from the very beginning. Logic programming languages, as well as other declarative languages viewed as specification tools found their early place between these two worlds. So did many mature formalisms, like RAISE [17], that incorporate both approaches.

One problem any specification method has to face is how exactly the general specification mechanisms used relate to the programs in a specific programming language. One solution to alleviate this, perhaps first put forward in Larch [18], is to extend explicitly the generic specification language by "interface languages" that link it to various programming languages. Another idea is to develop formalisms that interleave specification with programs, with examples like a number of languages devoted to *design by contract*, like Eiffel [24] and Extended ML [31, 21], or various extensions to programming languages to allow assertions to be inserted, which underlie a number of program verification systems and tools, like JML [9].

As far as applications are concerned, a recent survey on the use of formal methods [37] indicates that the specification/modeling activities were a crucial part of practically all reviewed projects that used any formal methods at all. No surprises here: even if the use of formal methods stops at providing a precise specification (necessarily involving some specification formalisms), the benefits of clear statement of the requirements on and properties of the system or component cannot be overestimated. In the early Transputer project, just the specification of the floating-point arithmetic [3] helped to discover problems with the standard to be implemented, and the full strength of the formal methods in use was highlighted even more when it came to formal verification of its Occam implementation and development of the chip microcode from it. Quite similar experience was brought for instance by a project to specify and verify the AAMP5 microprocessor in PVS (a specification language integrated with support tools and a theorem prover [28]). Two errors in the microcode were identified when the specification was given, and then a verification of the microcode was carried out [34].

These examples, as well as other examples (coming from different application areas, as for instance listed in [37]), show that this indeed is a typical pattern: a precise formal specification, often used to identify and clarify problems with informal design, offers a precise complete description of the systems (or their components) to be developed, a valuable asset on its own. Then formal methods and tools, often very specifically linked with the specification formalisms in use (with dependencies going in either direction) can be employed to support the development and verification of the software; see subsequent sections on model checking, verification, abstract interpretation, and the like. It is interesting that reports in [37] do not give a uniformly positive view of the benefits of the use of specification formalisms and related formal methods: while everybody agrees that their use results in an increased confidence and a higher quality of the product, their cost effectiveness is less clear. The major cost factor is, however, the high learning curve necessarily involved at the early stages of the project, when a (new) specification formalism is brought in and has to be learned to build the specifications; clearly, this overhead should decrease considerably with time when the use of a particular specification formalism is repeated and the use formal methods in general becomes an expected and required standard.

3.2 Verification Methods

We now describe briefly several major verification methods, abstract interpretation, model checking, equivalence checking, and verification by deduction. The properties to be proven are undecidable most of the time. In these cases, soundness—no false claim produced—and completeness—every property that holds is actually derived—can not be reached at the same time. The different methods deal with this in different ways. All are sound, will not derive correctness claims if, indeed, the system is faulty (*false positives*). They may be incomplete, producing warnings (*false negatives*) while, indeed, the system is correct, and/or they may require input from the user.

3.2.1 Abstract Interpretation

Abstract interpretation [12] computes an approximation of the program semantics. Alternatively, one might say that it computes *invariants*, properties that hold of every run of the system regardless of the environment. For example: At program line 17, variable `x` always has value 5, or the memory load at program point p is always a cache hit. As a signature we suggest

$$abstr\_int :: \mathsf{Sys} \to \mathsf{Spec}.$$

Abstract interpretation *per se* typically works on implementations without knowledge of input data and without specification. It infers properties that hold *for all possible executions*. In particular, let $abstr\_int(s) = \varphi$, and let $R = [\![\varphi]\!]$ be the set of runs denoted by φ. An abstract interpretation is called an *under-approximation*, if $R \subseteq [\![s]\!]$ and an *over-approximation* if $R \supseteq [\![s]\!]$. An abstract interpretation is always an over- or an under-approximation of the program semantics, $[\![s]\!]$. As an example, consider the safe classification of memory accesses as cache hits and cache misses as needed for a timing analysis of hard real-time programs described below. To classify memory accesses as cache hits one needs an under-approximation of the cache contents, while for cache misses one needs an over-approximation of the cache contents.

3.2.1.1 Derived Properties

Often, the specification (invariant) inferred by abstract interpretation is not the one one is actually interested in. Let φ_{spec} be the desired specification and let φ_{inv} be the inferred one. An abstract interpretation is *sound*, if $\varphi_{spec} \Rightarrow \varphi_{inv}$. In this case, one may obtain *false positives*, that is, indications that a specification is violated even though it is not. On the other hand, soundness allows the proof of absence of defects. If $\varphi_{inv} \Rightarrow \varphi_{spec}$ then the abstract interpretation is *complete*. If an error is found, it is definitely an error. Such a method is ideal for bug-hunting. On the other hand, a complete method might suffer from *false negatives*, that is, it might fail to uncover an error. While possible sound and complete abstract interpretations are possible, they are rare.

Often abstract interpretation is used with *implicit* specifications such as absence of bugs like division-by-zero, array-out-of-bounds-accesses, stack-overflow, null-pointer-dereferences. Typically, people deal with sound methods producing false positives rather than false negatives. On the other hand, abstract interpretation allows to *prove* the absence of certain defects making it attractive for certification of systems with respect to authority standards (e.g. in avionics).

3.2.1.2 Four Tools

Probably the largest industrial systems to which abstract interpretation was applied were safety-critical systems of the Airbus A380 plane. These systems consist of several hundred thousand lines of code. The four abstract interptretation based static analyzers described below, Polyspace Verifier, Astrée, Stackanalyzer, and aiT were able to analyze tasks of this size in times acceptable for the developers in the aeronautics and automotive industries.

Astrée is a sound static analyser for the programming language C designed to prove the absence of run-time errors such as division-by-zero, index-out-of-bounds, overflow and underflow, null, mis-aligned or dangling pointers [8]. Astrée was also designed to be complete, i.e., produce no false alarms, but only for the type of software found in critical real-time synchronous embedded control systems (e.g. synthesized from SCADE). Astrée[2] has been used with success in verifying aeronautic, aerospace, and automotive applications, such as electric flight control or space-vessels maneuvers, on programs up to 10^6 lines of code, without false alarms.

The sound static analysis tool with the largest user base is Polyspace Verifier[3]. It is in routine use in a decent set of development laboratories of the safety-critical embedded systems domain. Polyspace Verifier is based on abstract interpretation and analyzes programs written in C/C++ and Ada. Compared to Astrée it checks for absence of fewer errors and is less configurable. Its policy, "Green follows Orange" means that the analysis continues after a warning as if nothing happened. This means that several iterations are necessary to discover all problems.

Stackanalyzer[4] determines safe upper bounds on the size of system and user stacks. It determines the worst-case stack usage of the tasks in in the code under verification and displays the results as annotations of the call graph and control-flow graph.

Finally, aiT[5] determines safe upper bounds on the execution time of real-time programs [14]. Several different abstract interpretations are used, the most complex being the derivation of invariants about the set of all execution states of the execution platform. These invariants are used to bound the execution times of instructions. Depending on the complexity of the execution platform, aiT has shown an over-estimation of the execution times of between 8% for simple microcontrollers and 25% for complex high-performance mono-processors [35], while tasks of several million instructions can be analyzed within one day.

3.2.1.3 Roles in Abstract Interpretation

An abstract interpreter, as realized by a *tool developer* (based on foundations laid by the *researcher*) is able to analyze systems for a specific set of properties and nothing else. Any given tool is not universal, in contrast to underlying theory. Hopefully, the tool developer aims at the right set of properties for a relevant set of systems. The *user*, in principle, gets a push-button system. However, the analysis results may be much better if he/she gives a little help to the analyzer. This help may consist in configuring the system for the particular characteristics of the system, e.g. describing the ranges for environment variables and combining the right set of abstract domains for an embedded-control system in Astrée.

[2] http://www.absint.de/astree/
[3] http://www.mathworks.de/products/polyspace/
[4] http://www.absint.de/stackanalyzer/
[5] http://www.absint.de/ait/

It may also consist in supplying necessary properties about the execution platform, e.g. the type of used memory, or properties of the system under analysis, e.g. loop bounds, to aiT.

3.2.2 Model Checking

Model checking [11, 30] is understood by the following signature:

$$mc :: \mathsf{Model} \times \mathsf{Spec} \times \mathsf{Data} \to \mathbb{B} \times [\mathsf{Run}].$$

This means that given a model (or a system really) and a specification and input data, a model checker either provides a run to witness a possible error or indicates a successful check. Hence the signature has the "optional" result type [Run] with an extra bit representing satisfiable/unsatisfiable.

3.2.2.1 Model Checking in Industry

The most widespread use of model checking is in the semiconductor industry. Typical use cases were described by Cindy Eisner from IBM. A chip consists of several units. A *unit* is the smallest component of a processor architecture that has a functionality. A specification describing the functionality could be given for a unit. However, the state space to be exhaustively elaborated by a model checker currently is too large. Instead, *blocks*, parts of units, are checked. They may not have a specifiable functionality. So, only *local properties* are checked; for instance 14,000 local properties for the Pentium 4. These local properties express local correctness conditions.

Blocks have many different *environments* or *contexts*, in which they can be activated, in fact, too many to do this exhaustively. So, currently *blocks are checked for local properties in restricted sets of environments.*

Despite the dominance of model-checking use cases in hardware industry, there are examples from software and other industries as well. The *Static Driver Verifier Research Platform* [2] is a tool suite provided by Microsoft to verify Windows drivers. It is based on the software model checker SLAM.

3.2.2.2 Roles in Model Checking

Model checking has a different distribution of obligations from that of abstract interpretation. It places a higher burden on the user, who has to write a specification in the form of a finite-state machine or a temporal-logic formula, both not the native languages of most developers. In some cases, the user also has to supply an abstract model of the system, an often non-trivial task. This task may actually be alleviated by the recently begun cooperation of static analysts and model checkers. Abstraction of systems may be done based on the theory of abstract interpretation. The resulting abstract systems can then be model-checked.

3.2.3 Equivalence Checking

While model checking compares a model with its specification, equivalence checking compares two models. An adequate signature for equivalence checking also takes into account the provision of a counterexample in case that a difference between two systems is detected:

$$ec :: \mathsf{Spec}_1 \times \mathsf{Spec}_2 \to \mathbb{B} \times [\mathsf{Run}].$$

Spec_1 is typically referred to as the *Golden Implementation* and plays the role of a (formal) specification. Spec_2 is referred to as *Implementation*. Equivalence checking is mainly applied in the design of hardware systems, combinational or sequential circuits. The *Implementation* may be the result of adapting an existing design to a new semiconductor technology, performance optimization, compilation from register-transfer-level to gate-level, and the like. Very often, Spec_1 and Spec_2 share a lot of structural similarities, smoothing the way for the efficient application of graph-based data structures and algorithms (e.g., And-Inverter-Graphs or Binary Decision Diagrams) [22], as well as the application of SAT-solvers (for solving a satisfiability problem) [32] and ATPG-tools (for solving Automatic Test Pattern Generation problems), or combinations thereof [29].

3.2.4 Equivalence Checking in Industry

Industrial development processes of hardware designs routinely apply equivalence checking during the design process. Incremental and fine-grained design steps (as coarsely sketched in Fig. 1) ensure that the problem instances are manageable. Industrial applications of equivalence checking still require high level of user expertise for setting up the equivalence-checking framework, especially when implementations are delivered from outside customers or subcontractors.

3.2.5 Roles in Equivalence Checking

The roles in equivalence checking seem to be separated more clearly than for the other methods described in this section. The development team producing Spec_1 is typically different from the designers providing the implementation Spec_2. In the are of hardware development, the profession of a *Verification Engineer* was created. A verification engineer integrates Spec_1 and Spec_2 into the equivalence-checking framework while taking care of technological features, e.g., when some design features are deemed redundant with regard to the equivalence-checking task. Nevertheless, the verification engineer must be in close cooperation with the developers of the implementation, e.g., to get rid off *false-negative* counterexamples.

3.2.6 Verification by Deduction

Most abstractly, we describe verification by the signature

$$vbd :: \mathsf{Sys} \times \mathsf{Spec} \to \mathbb{B} \times [\mathsf{Proof}].$$

More precisely, the boolean result is either a proof that a system satisfies a given property; or a proof cannot be established. Logic is the lingua franca of *verification by deduction*. While logic is used in other approaches too (say model checking) it is really universal in verification. In interactive program verification, the user of a tool has to supply invariants at cutpoints of the program.

3.2.6.1 Academic and Industrial Practice

Verification of functional correctness by interactive theorem proving is standard practice for arithmetic units at processor manufacturers. This probably is the consequence of the Pentium bug [27, 13].

Proving the correctness of a compiler is, one could say, the "mother" of all software-verification attempts. The verification of the compiler guarantees that the safety properties

proved on the source code hold for the executable compiled code as well. Xavier Leroy developed and formally verified, i.e., gave a proof of semantic preservation, of a compiler from Clight (a large subset of the C programming language) to PowerPC assembly code. He used the Coq proof assistant both for programming the compiler and for proving its correctness [23].

C.A.R. Hoare's vision of the *Verifying Compiler* led to the Verified Software Repository (VSR). This is an evolving collection of tools and challenges related to software verification. It supports a community effort to develop technology to enable the mechanical certification of computer programs [5].

The Verisoft project [1] was a research project funded by the German Federal Ministry of Education and Research (BMBF). The main goal of the project was the pervasive formal verification of computer systems. The correct functionality of systems, as they are used, for example, in the automotive domain, in security technology and in medical technology, was to be mathematically proved. The proofs are computer aided in order to prevent human error by the scientists involved.

Finally, the verification effort developed within Microsoft Research cannot go unmentioned. Microsoft makes use of its verification platform Boogie[6]. Specifically, Boogie is an intermediate language generated by a number of front-end tools for specific purposes and languages like Havoc (pointer verification in C) or Chalice (concurrent). Boogie is probably the first example of information-exchange between verification engines. It has a wide selection of provers at its disposal to verify that programs adhere to their specs. Examples include Z3, Simplify, or Isabelle/HOL.

4 Acceptance

This section addresses the acceptance of formal methods in industry. It is based on the experience and observations of the participants, some industrial, some academic.

Compelling Needs

Different application domains have different requirements for verification and, therefore, call for the application of formal methods at different degrees.

- In general-purpose computing, time-to-market may be decisive, such that the additional cost of applying formal methods may be considered inappropriate. Users may be willing to tolerate system failures once in a while.
- For safety-critical embedded systems, failure is not acceptable. However, it is unrealistic to assume that complex systems consisting of millions of lines of code could be produced free of bugs. High-integrity subsystems still may be required to be free of bugs. The application of formal methods is mandatory to achieve the highest possible quality.
- For high-security systems or system components with high-security requirements, the existence of security loopholes is not tolerated. This area has the interesting feature that bug chasing is done by an external community, whether the designer want this or not, namely the hackers.

This classification is reflected in a recent survey of applications of formal methods in industry, see [7]. It lists a large number of applications in the transportation sector where

[6] http://research.microsoft.com/en-us/projects/boogie/

safety-critical subsystems have been developed using formal methods, followed by a good number of applications in hardware design and several in the financial sector.

Motivating the Introduction of Formal Methods

Often, a formal method is introduced into industrial practice after a major desaster that it could have prevented from occurring. A premier example is the Pentium bug. Formal verification of the Pentium's floating-point arithmetic unit could have saved Intel half a billion Dollars. This experience boosted the application of formal verification in the hardware industry. Similarly, several failures in medical instruments have cost the producers high liability costs and led to the introduction of formal methods.

The Role of Champions

Participants from industry emphasized the importance of champions in industry being enthusiastic about and competent in a formal method. Without these, a formal method is seldom introduced into industrial practice. On the other hand, an introduced method and tool may fall again into oblivion once a champion leaves his or her position.

The Role of Education

It is decisive for the acceptance of a method and tool in industry that the competences required from an industrial user are available or can be taught without too much effort. Industrial participants emphasized that teaching student how to develop high-quality code is more important than teaching them formal methods. This goes in the same direction as our emphasis of the importance of the design for verifyability. Disciplined, well-structured code will increase the applicability of formal methods and will allow for a higher degree of automation. At the same time, it is important that students obtain adequate mathematical background, enbaling them later to master the usage of formal-methods tools.

4.1 A Spectrum of Formality

The term *verification* is heavily abused, as was said above. Every activity expected to lead to a better system quality is subsumed under it. To account for this, we describe a spectrum of methods, not all considered "formal" that are used in practice.

Testing: Testing is still very popular despite C.A.R. Hoare's statement that it can only prove the existence and not the absence of bugs. In the terminology of Section 3.2, it is unsound and imcomplete.

Unsound static analysis: This method may be very helpful in chasing bugs [4]. However, it is neither sound nor complete and therefore not suitable for verification.

Model-based design: The most heavily used modeling languages have a brittle semantics, in fact, different semantics defined by different code generators. Only the ones having a formally defined semantics, e.g. Scade, would be subsumed under formal methods.

From lite to rigorous: Dines Bjørner in [7] describes a spectrum from lite, to rigorous, to formal application. "Lite" means means specifying the problem and maybe the solution in a formal specification language. "Rigorous" means to specify additional properties and possibly the relation between different specifications. "Formal" requires proofs of specified propositions.

5 Challenges and Perspectives

Industry representatives applying formal methods issued the following wish list, mainly concerning static-analysis tools.

- Improvements in tool functionality: higher precision, i.e., less false alarms, support for functionality analysis, better configurability, a way to trade precision for performance, stronger automation, diagnosis of the error–not of the error symptom– and possibly examples exhibiting the symptom.
- Improvements in scalability: Coping with very large systems: full verification for smaller programs, defect localization for large programs.
- Process support: compliance with chracteristics of the development and the certification processes, iterative and incremental verification, exploitation of model information avaiable in model-based design of safety-critical software, support for code quality assessments.
- Tool cooperation: support for information exchange between tools exploiting synergy between tools.

Keep it simple, predictable, actionable

Quoting from Tom Ball's presentation at the workshop, we understand that tools implementing formal methods are still too hard to use. Simple counterexamples and simple proofs are needed, as are predictable behavior to avoid wild swings in performance for small changes. Tools should explain their failures so that users know what to do next. Type systems are a particularly successful example satisfying these requirements. Finally, we should learn to build on each other's work and stop reinventing the wheel!

Design for Verifiability

A very recent area of research emerging from formal methods is design for verifiability. A number of design decisions influence the possibility and if this is given, the ease of verifying properties of systems. Traditionally, this is applied in programming-language design. The programming language may have a strong influence on the possibility and the needed effort of system analysis and verification. It may enforce restrictions whose validation would otherwise require an enormous effort. For example, a major problem in the analysis of imperative programs is the determination of dependences between the statements in programs. The unrestricted use of pointers, as in the C programming language, makes this analysis of dependences very difficult due to the severe alias-problem created through pointers.

Several coding guidelines have been proposed to lead to more disciplined code. They typically restrict the use of the *dark corners* of the programming language, e.g. pointer arithmetic and function pointers. One prominent example is MISRA C [25], the C coding standard proposed by the Motor Industry Software Reliability Association. It bans the worst features of C with respect to the software verification task. However, this does not necessarily lead to programs whose timing behavior is precisely predictable [15].

The execution platform determines the analyzability of the timing behavior [36]. Most emerging multi-core platforms will make timing analysis infeasible due to the interference of threads on shared resources.

The transition from federated architectures—one computer per function—to integrated architectures—several functions integrated on one platform—as currently under way in the

Integrated Modular Avionics (IMA) and the AUTomotive Open System Architecture (AUTOSAR) standardization efforts offers a great chance to improve the verifiability of systems. Temporal and spatial partitioning is used in IMA to avoid the logical interference of functions. However, the existing implementations give away the chance to cleanly and efficiently deal with the interaction on shared resources and the resulting non-composability of the resource behavior. In the ideal case, *design meets verification*, that is, design only admits systems that can be easily verified.

Limitations

However close we get to modelling real systems, we will always talk about models abstracting from some details. The same goes for specifications. Details left unmodeled and/or unspecified cannot be verified, obviously, and remain a fundamental limitation.

Acknowledgements

We thank Marc Herbstritt for supporting us during the meeting by taking minutes which built a basis for this manifesto and for providing a draft on equivalence checking.

6 Participants

- Krzysztof Apt
CWI – Amsterdam, NL
- Thomas Ball
Microsoft Res. – Redmond, US
- Dines Bjørner
Holte, DK
- Patrick Cousot
ENS – Paris, FR
- Cindy Eisner
IBM – Haifa, IL
- Javier Esparza
TU München, DE
- Steffen Görzig
Daimler AG – Böblingen, DE
- Yuri Gurevich
Microsoft Res. – Redmond, US
- Marc Herbstritt
Schloss Dagstuhl, DE
- Manuel Hermenegildo
IMDEA Software – Madrid, ES
- Bengt Jonsson
University of Uppsala, SE
- Joseph Roland Kiniry
IT Univ. of Copenhagen, DK
- Jörg Kreiker
TU München, DE
- Wei Li
Beihang University, CN
- Wolfgang J. Paul
Universität des Saarlandes, DE
- Erik Poll
Radboud Univ. Nijmegen, NL
- Sriram K. Rajamani
Microsoft Research India – Bangalore, IN
- Jean-Francois Raskin
Univ. Libre de Bruxelles, BE
- John Rushby
SRI – Menlo Park, US
- Donald Sannella
University of Edinburgh, GB
- Wei Sun
Beihang University, CN
- Andrzej Tarlecki
University of Warsaw, PL
- Wolfgang Thomas
RWTH Aachen, DE
- Moshe Y. Vardi
Rice University, US
- Reinhard Wilhelm
Universität des Saarlandes, DE
- Jim C.P. Woodcock
University of York, GB
- Lenore Zuck
NSF – Arlington, US

References

1 Eyad Alkassar, Mark A. Hillebrand, Dirk Leinenbach, Norbert W. Schirmer, and Artem Starostin. The Verisoft approach to systems verification. In Natarajan Shankar and Jim Woodcock, editors, *Verified Software: Theories, Tools, Experiments Second International Conference, VSTTE 2008, Toronto, Canada, October 6–9, 2008. Proceedings*, volume 5295 of *Lecture Notes in Computer Science*, pages 209–224, Toronto, Canada, October 2008. Springer.

2 Thomas Ball, Ella Bounimova, Vladimir Levin, Rahul Kumar, and Jakob Lichtenberg. The static driver verifier research platform. In Tayssir Touili, Byron Cook, and Paul Jackson, editors, *CAV*, volume 6174 of *Lecture Notes in Computer Science*, pages 119–122. Springer, 2010.

3 Geoff Barrett. Formal methods applied to a floating-point number system. *IEEE Transactions on Software Engineering*, 15(5):611–621, 1989.

4 Al Bessey, Ken Block, Benjamin Chelf, Andy Chou, Bryan Fulton, Seth Hallem, Charles Henri-Gros, Asya Kamsky, Scott McPeak, and Dawson R. Engler. A few billion lines of code later: using static analysis to find bugs in the real world. *Commun. ACM*, 53(2):66–75, 2010.

5 Juan Bicarregui, C. A. R. Hoare, and J. C. P. Woodcock. The verified software repository: a step towards the verifying compiler. *Formal Asp. Comput.*, 18(2):143–151, 2006.

6 Michel Bidoit and Peter D. Mosses, editors. *CASL User Manual*. Number 2900 in Lecture Notes in Computer Science. Springer, 2004.

7 Dines Bjøner. A spring 2011 survey of applications of formal methods in industry, a draft version. http://www2.imm.dtu.dk/ db/fm-industry-project-survey/fm-applics-survey.pdf, 2011.

8 Bruno Blanchet, Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, David Monniaux, and Xavier Rival. A static analyzer for large safety-critical software. In *PLDI*, pages 196–207. ACM, 2003.

9 L. Burdy, Y. Cheon, D.C. Cok, M.R. Ernst, J.R. Kiniry, G.T. Leavens, K.R.M. Leino, and E. Poll. An overview of JML tools and applications. *International Journal on Software Tools for Technology Transfer (STTT)*, 7(3):212–232, 2005.

10 R.M. Burstall and J.A. Goguen. An informal introduction to specifications using Clear. In R.S. Boyer and J.S. Moore, editors, *The Correctness Problem in Computer Science*, pages 185–213. Academic Press, 1981. Also in: *Software Specification Techniques* (eds. N. Gehani and A.D. McGettrick), Addison-Wesley, 1986.

11 Edmund M. Clarke, E. Allen Emerson, and A. Prasad Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Trans. Program. Lang. Syst.*, 8(2):244–263, 1986.

12 Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252, 1977.

13 Alan Edelman. The mathematics of the Pentium division bug. *SIAM Rev.*, 39:54–67, 1997.

14 Christian Ferdinand, Reinhold Heckmann, Marc Langenbach, Florian Martin, Michael Schmidt, Henrik Theiling, Stephan Thesing, and Reinhard Wilhelm. Reliable and precise WCET determination for a real-life processor. In *EMSOFT*, volume 2211 of *LNCS*, pages 469–485. Springer, 2001.

15 Gernot Gebhard, Christoph Cullmann, and Reinhold Heckmann. Software structure and WCET predictability. In *Proc. of the workshop Bringing Theory to Practice: Predictability and Performance in Embedded Systems*, Grenoble, 2011.

16 Joseph Goguen, James Thatcher, and Eric Wagner. An initial algebra approach to the specification, correctness and implementation of abstract data types. Technical Report RC 6487,

IBM Watson Research Center, Yorktown Heights NY, 1976. Also in: *Current Trends in Programming Methodology. Volume IV (Data Structuring)* (ed. R.T. Yeh), Prentice-Hall, 80–149, 1978.

17 The RAISE Language Group. *The RAISE Specification Language*. The BCS Practitioners Series. Prentice-Hall, 1992.

18 John V. Guttag and James J. Horning. *Larch: Languages and Tools for Formal Specification*. Springer, 1993.

19 C.A.R. Hoare. An axiomatic basis for computer programming. *Communications of the Association for Computing Machinery*, 12:576–580,583, 1969.

20 Cliff B. Jones. *Software Development: A Rigorous Approach*. Prentice-Hall, 1980.

21 Stefan Kahrs, Donald Sannella, and Andrzej Tarlecki. The definition of Extended ML: A gentle introduction. *Theoretical Computer Science*, 173:445–484, 1997.

22 Andreas Kuehlmann, Viresh Paruthi, Florian Krohm, and Malay K. Ganai. Robust boolean reasoning for equivalence checking and functional property verification. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 21(12):1377–1394, 2002.

23 Xavier Leroy. Formal verification of a realistic compiler. *Communications of the ACM*, 52(7):107–115, 2009.

24 Bertrand Meyer. *Eiffel : The Language*. Prentice-Hall, 1991.

25 MISRA. *Guidelines for the Use of the C Language in Critical Systems*, October 2004. ISBN 0 9524156 2 3.

26 Peter D. Mosses, editor. *CASL Reference Manual*. Number 2960 in Lecture Notes in Computer Science. Springer, 2004.

27 Thomas R. Nicely. The Pentium division flaw. *Virginia Scientists Newsletter*, 1:3, 1995.

28 Sam Owre, John M. Rushby, and Natarajan Shankar. PVS: A prototype verification system. In Deepak Kapur, editor, *Automated Deduction — CADE-11, 11th International Conference on Automated Deduction, Saratoga Springs, NY, USA, June 15-18, 1992, Proceedings*, volume 607 of *LNCS*, pages 748–752. Springer, 1992.

29 Viresh Paruthi and Andreas Kuehlmann. Equivalence checking combining a structural sat-solver, bdds, and simulation. In *Proc. IEEE Int'l Conference On Computer Design (ICCD), Austin, USA*, pages 459–464, 2000.

30 Jean-Pierre Queille and Joseph Sifakis. Specification and verification of concurrent systems in Cesar. In *Symposium on Programming*, volume 137 of *Lecture Notes in Computer Science*, pages 337–351. Springer, 1982.

31 Donald Sannella and Andrzej Tarlecki. Program specification and development in Standard ML. In *Proceedings of the 12th ACM Symposium on Principles of Programming Languages*, pages 67–77, 1985.

32 João P. Marques Silva and Karem A. Sakallah. Boolean satisfiability in electronic design automation. In *DAC*, pages 675–680, 2000.

33 J. Michael Spivey. *The Z Notation: A Reference Manual*. Prentice Hall International Series in Computer Science, second edition, 1992.

34 Mandayam K. Srivas and Steven P. Miller. Applying formal verification to the AAMP5 microprocessor: A case study in the industrial use of formal methods. *Formal Methods in System Design*, 8(2):153–188, 1996.

35 Lili Tan. The worst-case execution time tool challenge 2006. *STTT*, 11(2):133–152, 2009.

36 Lothar Thiele and Reinhard Wilhelm. Design for timing predictability. *Real-Time Systems*, 28(2-3):157–177, 2004.

37 Jim Woodcock, Peter Gorm Larsen, Juan Bicarregui, and John S. Fitzgerald. Formal methods: Practice and experience. *ACM Computing Surveys*, 41(4), 2009.

Manifesto from Dagstuhl Perspectives Workshop 11331

Improving The Future of Research Communications and e-Scholarship

Edited by

Philip E. Bourne[1], Timothy W. Clark[2], Robert Dale[3], Anita de Waard[4], Ivan Herman[5], Eduard H. Hovy[6], and David Shotton[7]

1 **University of California at San Diego, US,** pbourne@ucsd.edu
2 **Harvard Medical School, US,** tim_clark@harvard.edu
3 **Macquarie University, AU,** Robert.Dale@mq.edu.au
4 **Elsevier Laboratories – Jericho, US,** A.dewaard@elsevier.com
5 **Centrum voor Wiskunde en Informatica – Amsterdam, NL,** ivan@w3.org
6 **University of Southern California – Marina del Rey, US,** hovy@isi.edu
7 **University of Oxford, GB,** david.shotton@zoo.ox.ac.uk

Abstract

The dissemination of knowledge derived from research and scholarship has a fundamental impact on the ways in which society develops and progresses, and at the same time it feeds back to improve subsequent research and scholarship. Here, as in so many other areas of human activity, the internet is changing the way things work; two decades of emergent and increasingly pervasive information technology have demonstrated the potential for far more effective scholarly communication. But the use of this technology remains limited.

Force11 is a community of scholars, librarians, archivists, publishers and research funders that has arisen organically to help facilitate the change toward improved knowledge creation and sharing. This document highlights the findings of the Force11 workshop on the Future of Research Communication held at Schloss Dagstuhl, Germany, in August 2011: it summarizes a number of key problems facing scholarly publishing today, and presents a vision that addresses these problems, proposing concrete steps that key stakeholders can take to improve the state of scholarly publishing.

Seminar 15.–18. August, 2011 – www.dagstuhl.de/11331
1998 ACM Subject Classification K. Computing Milieux
Keywords and phrases Science publishing, online communities, science policy, digital repositories, semantic publishing, citation analysis, data publication information access and integration, reporting standards
Digital Object Identifier 10.4230/DagMan.1.1.41

Executive Summary

Research and scholarship lead to the generation of new knowledge. The dissemination of this knowledge has a fundamental impact on the ways in which society develops and progresses, and at the same time it feeds back to improve subsequent research and scholarship. Here, as in so many other areas of human activity, the internet is changing the way things work: it opens up opportunities for new processes that can accelerate the growth of knowledge, including the creation of new means of communicating that knowledge among researchers and within the wider community. Two decades of emergent and increasingly pervasive information technology have demonstrated the potential for far more effective scholarly communication. However, the use of this technology remains limited; research processes and the dissemination of research results have yet to fully assimilate the capabilities of the web and other digital media. Producers and consumers remain wedded to formats developed in the era of print publication, and the reward systems for researchers remain tied to those delivery mechanisms.

Force11 (the Future of Research Communication and e-Scholarship) is a community of scholars, librarians, archivists, publishers and research funders that has arisen organically to help facilitate the change toward improved knowledge creation and sharing. Individually and collectively, we aim to bring about a change in scholarly communication through the effective use of information technology. Force11 has grown from a small group of like-minded individuals into an open movement with clearly identified stakeholders associated with emerging technologies, policies, funding mechanisms and business models. While not disputing the expressive power of the written word to communicate complex ideas, our foundational assumption is that scholarly communication by means of semantically-enhanced media-rich digital publishing is likely to have a greater impact than communication in traditional print media or electronic facsimiles of printed works. However, to date, online versions of 'scholarly outputs' have tended to replicate print forms, rather than exploit the additional functionalities afforded by the digital terrain. We believe that digital publishing of enhanced papers will enable more effective scholarly communication, which will also broaden to include, for example, better links to data, the publication of software tools, mathematical models, protocols and workflows, and research communication by means of social media channels.

This document highlights the findings of the Force11 Dagstuhl Perspectives Workshop on "The Future of Research Communication" held at Schloss Dagstuhl, Germany, in August 2011: it summarizes a number of key problems facing scholarly publishing today, and presents a vision that addresses these problems, proposing concrete steps that key stakeholders can take to improve the state of scholarly publishing. More about Force11 can be found at http://www.force11.org [16]. This manifesto is a collaborative effort that reflects the input of all Force11 attendees at the Dagstuhl Perspectives Workshop (see Sect. 7).[1]. We see it as a starting point that will grow and be updated and augmented by individual and collective efforts by the participants and others. We invite you to join and contribute to this enterprise.

[1] See also the corresponding Force11 White Paper, which is very much a living document: Bourne P, Clark T, Dale R, de Waard A, Herman I, Hovy E and Shotton D (eds.), on behalf of the Force11 community (2011). Force11 White Paper:

About This Document

This document contains five sections. Section 1 presents our vision of the future of scholarly publishing. In Section 2, we outline six key problems that prevent scholarly communication from achieving its full potential. Section 3 contains six specific recommendations for actions to address these problems. Section 4 offers a dynamic list of pointers to relevant research reports and related projects. Finally, in Section 5 we describe what we are doing to implement these recommendations.

The problems and recommendations we perceive can be grouped into two groups, each containing three principal themes:

- Themes 1–3 concern the *format and technologies* of scholarly publication: how scholarly data, information, and knowledge are (or could be) represented; how readers, users, authors, editors and computers can interact with these representations; and how different knowledge representations could be combined, queried, stored and otherwise treated.
- Themes 4–6 concern the enterprise of scholarly publishing, including business models and the attribution of credit. In these sections we discuss how scholarship is evaluated, accredited and monetized; current and new models and modes of assigning copyright and intellectual property rights; the financial aspects of scholarly publishing; and the mechanisms for assessing the quality and value of researchers and their research outputs, and of attributing credit and worth to them.

The problems relating to these six themes are described in Section 2, while our recommendations for their solutions are described in Section 3. These problems and recommendations are summarized in the following table.

Problems	Recommendations
Formats and Technologies	
2.1 Existing formats needlessly limit, inhibit and undermine effective knowledge transfer	3.1 Rethink the unit and form of the scholarly publication
2.2 Improved knowledge dissemination mechanisms produce information overload	3.2 Develop tools and technologies that better support the scholarly lifecycle
2.3 Claims are hard to verify and results are hard to reuse	3.3 Add data, software, and workflows into the publication as first-class research objects
Business Models and Attribution of Credit	
2.4 There is a tension between commercial publishing and the provision of unfettered access to scholarly information	3.4 Derive new financially sustainable models of open access
2.5 Traditional business models of publishing are being threatened	3.5 Derive new business models for science publishers and libraries
2.6 Current academic assessment models don't adequately measure the merit of scholars and their work over the full breadth of their research outputs	3.6 Derive new methods and metrics for evaluating quality and impact that extend beyond traditional print outputs to embrace the new technologies

Table of Contents

1 Our Vision

A dispassionate observer, perhaps visiting from another planet, would surely be dumbfounded by how, in an age of multimedia, smartphones, 3D television and 24/7 social network connectivity, scholars and researchers continue to communicate their thoughts and research results primarily by means of the selective distribution of ink on paper, or at best via electronic facsimiles of the same.

Modern technologies enable vastly improved knowledge transfer and far wider impact. Freed from the restrictions of paper, numerous advantages appear. Communication becomes instantaneous across geographic boundaries. Terms in electronic documents may be automatically disambiguated and semantically defined by linking to standard terminology repositories, allowing more accurate retrieval in searches; complex entities mentioned in documents may be automatically expanded to show diagrams or pictures that facilitate understanding; citations to other documents may be enhanced by summaries generated automatically from the cited documents. Documents may be automatically clustered with others that are similar, showing their relationship to others within their scholarly context, and their place in the ongoing evolution of ideas. Ancillary material that augments the text of the scholarly work may be linked to or distributed with the work; this may include numerical data (from experiments), images and videos (showing procedures or scenarios), sound recordings, presentational materials, and other elements in forms of media still on the horizon. Extracts and discussions of scholarly work on social media such as blogs, online discussion groups and Twitter may greatly broaden the visibility of a work and enable it to be better evaluated and cross-linked to other information sources. A broad range of recent technological advances provide increasingly diverse and powerful opportunities for more effective scholarly communication; we need to grasp the opportunities and make these possibilities realities.

We see a future in which scientific information and scholarly communication more generally become part of a global, universal and explicit network of knowledge; where every claim, hypothesis, argument—every significant element of the discourse—can be explicitly represented, along with supporting data, software, workflows, multimedia, external commentary, and information about provenance. In this world of networked knowledge objects, it would be clear how the entities and discourse components are related to each other, including relationships to previous scholarship; learning about a new topic means absorbing networks of information, not individually reading thousands of documents. Adding new elements of scholarly knowledge is achieved by adding nodes and relationships to this network. People could contribute to the network from a variety of perspectives; each contribution would be immediately accessible globally by others. Reviewing procedures, as well as reputation management mechanisms, would provide ways to evaluate and filter information. This vision moves away from the Gutenberg paper-centric model of the scholarly literature, towards a more distributed network-centric model; it is a model far better suited for making knowledge-level claims and supporting digital services, including more effective tracking and interrogation of what is known, not known, or contested.

To enable this vision, we need to create and use new forms of scholarly publication that work with reusable scholarly artifacts. Two principal aspects can be distinguished. First, we need to revise the artifacts of communication. As a starting point, our vision entails creating a new, enriched form of scholarly publication that enables the creation and management of relationships between knowledge, claims and data. It also means the creation of a knowledge infrastructure that allows the sharing of computationally executable

components, such as workflows, computer code and statistical calculations, as scientifically valid content components; and an infrastructure that allows these components to be made accessible, reviewed, referenced and attributed. To do this, we have to develop best practices for depositing research datasets in repositories that enable linking to relevant documents, and that have high compliance levels driven by appropriate incentives, resources and policies. In addition, for scientific domains, the new forms of publication must facilitate reproducibility of results, which means, at least for *in silico* research, the ability to preserve and re-perform executable workflows or services. This will require the ability to re-construct the context in which these objects were executed, which may well contain or reference other executable objects as well as data objects that may evolve through time. In this way, the content of communications about research will follow the same evolutionary path that we have seen for general web content: a move from the static to the increasingly dynamic.

With all this, we do recognise the importance of the peer-reviewed journal article as a primary dissemination channel and public record of new research results, since it uniquely provides a dated version of record of the authors' views at the time of publication, and as such becomes an immutable part of the scientific record. But even here, with this the most traditional of scholarly communication media, we can with existing technologies provide immediate improvements: semantic enhancements to the text; greater interactivity with tables and figures; access to the data within articles in actionable form; data fusions (mashups) with data from other sources, for example Google Maps, where appropriate; direct citation of and links to underlying datasets stored in databases and data repositories; and the open publication in machine-readable form of both the full bibliographic record for the article and also the citation information contained within the article's reference list, encoded using appropriate ontologies, so that these basic facts can enter the web of linked open data [36, 31].

The second component of our vision requires changes to the complex socio-technical scholarly and commercial ecosystem. In particular, to obtain the benefits that networked knowledge promises, we have to put in place reward systems that encourage scholars and researchers to participate and contribute. We need to acknowledge the fact that notions such as journal impact factor are poor surrogates for measuring the true impact of scholarship, and are increasingly irrelevant in a world of disaggregated knowledge units of vastly varying granularity; and we need to derive new mechanisms that allow us more accurately to measure true contributions to the ongoing enterprise of augmenting the world's store of knowledge. The business models that are currently driving scholarly publishing, which rest mainly on libraries buying access rights to digital journals from publishers, are clearly no longer adequate to support the rich, variegated, integrated and disparate knowledge offerings that new technologies enable, and that new scholarship requires. In a collaboration involving scholars, publishers, libraries, funding agencies, and academic institutions, we need to develop models that can enable this exciting future to develop, while offering sustainable forms of existence for the constituent parties, although perhaps not in their present states.

If we get this right, the potential is immense. The changes we envisage pave the way for a revolution in the manner in which research is carried out and communicated, leading to significant improvements in scholarly productivity and quality, and enhanced transparency that can only increase the public's trust in the value of science. Similar benefits apply to scholarship in the arts and humanities.

These developments bring advantages for many parties:

- **For scholars** (also in their roles as authors, editors and reviewers) the benefits are better communication of knowledge: easier transmission of information from its creators or discovers (the producers), in more forms using richer media, permitting easier, faster

and deeper interpretation of the information by the consumers (other scholars, students and their teachers, government and non-governmental agencies, industry, the media, and society at large). At the same time, these new and enhanced forms of communication will enable more accurate evaluations of the quality and the impact of scholars' work, facilitating better promotion evaluations and proposal assessments.

- Similarly, **for decision makers and managers**, the new communicative forms mean that the impacts and effects of scholarly communications, and hence of their authors, can more easily be tracked and evaluated.
- For **research funders**, enhanced communications will enable more accurate overviews of the size, direction and importance of each stream of research, and permit quicker determination of the quality of the work cited in grant proposals. But these advances mean that established practice will need to change.
- For **librarians and archivists**, while online accessibility will mean that traditional library holdings become less important, the archiving, updating and maintenance of digital data and software will increase in importance. Adapting to these changes will bring about new modes of service to users.
- Similarly, **for publishers**, the traditional functions of manuscript compilation and distribution will change radically, while quality control, access facilitation, new modes of aggregation, and the standardization, maintenance, and support of knowledge access technologies become more important. Providing these services will allow publishers successfully to face the challenges of free access to published research that is being ushered in by the open access movement.

2 Problem: The Growing Problems of Outdated Communication

We are a long way from achieving this vision today. As noted above, the impediments exist primarily in two dimensions: we have to change the *nature of the formats and technologies of communication*, that underpins the world of scholarly publishing, and we have to change the *social ecosystem of communication* that has grown up around the existing technologies. We review the key issues in these two areas in turn.

Problems with Current Formats and Technologies

2.1 Existing Formats Are Not Tailored for Knowledge Transfer

Scholarly communications are, at this mid-point in the digital revolution, in an ill-defined transitional state—a 'horseless carriage' state—that lies somewhere between the world of print and paper and the world of the web and computers, with the former still exercising significantly more influence than the latter. However, the recent development of new media and communicative possibilities using information technology, and the need to communicate and comprehend increasing amounts of additional information such as numerical and multimedia data, make the traditional forms inadequate. Continued reliance on paper documents and their electronic shadows make it very difficult or impossible to incorporate massive amounts of data, moving images or software; there is simply no natural way to associate such ancillary information 'into' the traditional publication. Additionally, any software-based text mining

or information extraction procedures require that paper-based information first be converted into machine-tractable form and made freely available for such mining.

2.2 The Ever-Increasing Problem of Information Overload

Scholars have experienced information overload for more than a century [35] and the problem is just getting worse. Online access provides much better knowledge discovery and aggregation tools, but these tools struggle with the fragmentation of research communication caused by the rapid proliferation of increasingly specialized and overlapping journals, some with decreasing quality of reviewing [29].

2.3 Verifying Claims and Re-using Results

Most types of scholarship involve claims, and all sciences and many other fields require that these claims be independently testable. Good results are often re-used, sometimes thousands of times. But actually obtaining the necessary materials, data or software for such re-use is far harder than it should be. Even in the rare cases where the data are part of the research communication, these are typically relegated to the status of 'supplementary material', whose format [23] and preservation [27] are inadequate. Sometimes the data are archived in separate data repositories that offer a more secure long-term future. But in such circumstances efforts need to be made to ensure that their links to the relevant textual research communications are explicit, robust and persistent. At present it is difficult for a scholar easily and sustainably to record the data on which the work is based in a form that others can absorb and use, and to maintain links to the associated textual publication.

Problems With Business and Assessment Models

2.4 Next-generation Tools Require Unfettered Resource Access

Currently, a large and active movement of professionals and students, including data curators, are providing services intended to improve the effectiveness of scholarly communication, and thereby the productivity of researchers; these entail digging facts out of textual publications and presenting them in machine-readable actionable form. The need for much of this expensive manual effort would be reduced if authors were to provide the relevant metadata at the time of publication. These extraction processes are increasingly being performed by automated text mining and classification software. However, because the source material is usually copyrighted, and these rights are distributed across a large number of publishers, the service providers are forced to negotiate individual contracts with each publisher, which is extremely wasteful of time and resources. To reduce this burden, some research funders are increasingly mandating that research results of all types be made openly available. However, this results in a confusing world where some publications are immediately and freely available and others on the same topic are not.

A related problem is the effect of the web as the medium for scholarly communication, since it is ending the role of local library collections. Libraries and archives have been forced to switch from purchasing copies of the research communications of interest to their readers, to leasing web access to the publishers' copies, with no assurance of long-term accessibility

to current content if future subscriptions lapse. Bereft of almost all their original value to scholars, libraries are being encouraged to both compete in the electronic publishing market and to take on the task of running 'institutional repositories', in effect publishing their scholars' data and research communications. Though both tasks are important, neither has an attractive business model. Re-publishing an open access version of their scholars' output where research is published in subscription-access journals may seem redundant, but it is essential if the artificial barriers that intellectual property restrictions have erected to data-mining and other forms of automated processing are to be overcome [20].

2.5 Traditional Publishing Models Are Under Attack

Academic publishers have been slower to encounter, but are not immune from, the disruption that the internet has wrought on other content industries [34]. The academic publishers' major customers, academic libraries, are facing massive budget cuts [22], and so are unlikely to be a major source of continued revenue. The internet has greatly reduced the costs of publishing, new players (such Google and other software companies) have appeared in the market, and legislative and funding bodies are actively addressing issues of free access to data and text [20]. The advent of the internet has greatly reduced the monetary value that can be extracted from paper-based academic content, and science publishers, who have traditionally depended on extracting this value, face a crisis, since their old business models are suffering disruption. Conversely, the internet permits the creation of new added-value services relating to search, semantics and integration that present exciting new commercial opportunities. Clearly the scholarly publishing industry needs to engage in discussions with different partners within the value chain, if it is to be included in the development of the new standards, services, business models, metrics/analysis, legislation, knowledge ecosystems and evaluation frameworks that the internet now makes possible, rather than being supplanted by new agile startups that have the ability to adapt more swiftly.

The software developers who build the current research informatics infrastructure are also very aware of the shortfalls and hindrances generated by today's fragmented development efforts. The problems here can be attributed to a number of elements. First, heterogeneous technologies and designs, and the lack (or sometimes the superfluity!) of standards, cause unnecessary technical difficulties and directly affect integration costs. Second, a complex landscape of intellectual property rights and licensing for software add legal concerns to developers' requirements. Third, research software developers typically work in a competitive environment, either academic or commercial, where innovation is rewarded much more highly than evolutionary and collaborative software reuse. This is especially true in a funding environment driven by the need for intensive innovation, where reusing other peoples' code is a likely source of criticism. Finally, even under optimal technical conditions, it is still challenging for software programmers to understand what components are the most appropriate for a given challenge, to make contact with the correct people to facilitate the construction of tools, and to work within distributed teams across groups to build high-quality interoperable software. The impact of these tools is, far too often, solely based on how immediately useful they will be to researchers themselves, with no thought for the wider community.

Thus changing roles and business models form an immense challenge for libraries, publishers and software developers. The only fruitful way forward, we firmly believe, will be for all parties collaborating to build new tools that optimally support scholarship in a distributed

open environment. Only by creating a demonstrably better research environment will we convince the entire system of scholarly communication and merit assessment to adopt new forms and models.

2.6 Current Assessment Models Don't Measure Merit

Not only are the products of research activity still firmly rooted in the past, so too are our means of assessing the impact of those products and of the scholars who produce them. For five decades, the impact of a scholarly work—an entity that is already narrowly defined, in the sciences as a journal article, and in the humanities as a monograph—has been judged by counting the number of citations it receives from other scholarly works, or, worse, by attributing worth to an individual's work based solely on the overall impact factor of the journal in which it happens to be published. We now live in an age in which other methods of evaluation, including article-level usage metrics, blog comments, discussion on mail lists, press quotes, and other forms of media, are becoming increasingly important reflections of scholarly and public impact. Failure to take these aspects into account means not only that the impact and/or quality of a publication is not adequately measured, but also that the current incentivization and evaluation system for scholars does not relate well to the actual impact of their activities.

3 Strategies for Change

Mirroring our identification of the six impediments to our vision that lie in the two dimensions of technology and society, we here make specific recommendations for change in these two dimensions.

New Publication Formats and Tools

3.1 Rethink the unit and form of the scholarly publication: the Research Object

At the foundation of any change is the infrastructure to support that change. One must no longer think of the journal article or research paper as the standard unit of currency by which knowledge is exchanged. Now it is but one among many forms. In the most generic sense, the new form of knowledge exchange centers on the *research object* [11, 5], a container for a number of related digital objects—for example a paper with associated datasets, workflows, software packages, etc., that are all the products of a research investigation and that together encapsulate some new understanding. Publishing of research objects is not necessarily publishing as we know it today, achieved by the same mechanisms as used for traditional scholarly articles. It consists of providing free and open access to the component parts of the research object, that may or may not have been individually reviewed by others either pre- or post-publication.

Arriving at a suitable definition of research objects requires work on standards and provenance, and conformance to general principles, some of which are suggested here:

- Support for multiple media types—text, images, podcasts, videos, etc.

- Recognition that raw and derived data, data processing procedures, computational models, experimental protocols and workflows all need to be preserved as part of the research object, and shared publicly.
- Support for access to content at varying granularities of detail.
- Support for the automatic extraction of information from research objects at these varying granularities, and its integration with third-party information.
- Support for uniquely identifying all elements of the research object.
- Support for both human and machine access, including access by disabled humans.
- Support for existing and emerging web and semantic web standards surrounding data representation and linking.
- Inclusion of social media as legitimate components within the world of the scientific discourse.

The research object per se does not necessarily capture the processes by which research leads to new knowledge. There is a temporal aspect to research and the scholarly lifecycle that also needs to be recorded, either within research objects or between research objects, and that should also be capable of being reproduced.

Developing the tools to support these changes, if undertaken from scratch, would be an immense undertaking. Thus, where possible, existing tools should be adapted and integrated within the new open infrastructure. Several classes of tools that exist and could be considered as components for this infrastructure are detailed in the "Tools" section of the Force11 web site [16].

What is happening now?

The following are examples of technological changes associated with new forms of scholarship.[2]

- Hypothesis/claim-based representation of the rhetorical structure of a scientific paper [13].
- Modular formats for science publishing [14].
- Developments of metadata standards and ontologies for describing publishing activities and publications, for characterizing citations between them, for identifying their structural and rhetorical components, and for describing discourse elements within the text.
- Semantic publishing initiatives and other enriched forms of publication.

What are the next steps?

Change is likely to occur gradually through a series of incremental steps, most of which will not be driven by the technology. Rather, the technology should respond to the recognized requirements of scientists for improved dissemination, reproducibility, recognition, etc. These requirements need to be assessed and formalized. The very existence of Force11 is an acknowledgement of the need for changes, but these changes need to be quantified and specifications drawn up for their solution.

[2] Readers should also consult our online collection of links to related activities and examples at https://sites.google.com/site/futureofresearchcommunications/links/links.

3.2 Develop tools and technologies that better support the scholarly lifecycle

What is happening now?

As scholarship in all fields increasingly becomes undertaken online, new tools and technologies are required to support the whole scholarly lifecycle from initial hypothesis to results publication. We are already seeing:

- the emergence of workflow systems;
- the emergence of data repositories within which datasets have globally unique identifiers and explicit links to journal articles, which by necessity provide some form of attribution and provenance information;
- the emergence of citation ontologies and corpora of open citation data;
- the emergence of software repositories with good versioning support; and
- the increasing use of online services for collaborative work: file exchange services such as Dropbox, collaborative note-taking environments such as EtherPad, and collaborative authoring environments such as Google Docs.

Nevertheless, these systems are acknowledged to be inadequate and cumbersome in their use. We require:

- better systems to permit collaborative work by geographically distributed colleagues;
- better systems to permit collaborative writing, with fail-safe versioning;
- better tools for richer interactive data and metadata visualization, enabling dynamic exploration; and
- easier data publication mechanisms, including better integration with data acquisition instrumentation, so that the process becomes automated.

What are the next steps?

To begin with, we want the scholarly community to be concerned with modes of archiving and sharing papers, data, workflows, models and software, and with the creation of research objects as part of their daily research routines. Other questions to explore include:

- What are the features of the research lifecycle and how do they impact the contents of and relationships between the artefacts that constitute digital research objects?
- How can existing tools be adapted to fit the specific workflow requirements of different scholarly domains?
- How can these tools be optimally integrated with environments to read, write and edit publications, and to create and evaluate research data?

3.3 Integration of datasets, software, mathematical models and workflows into publications as first-class research objects

Clearly, data in 21st Century science are almost always subjected to transformation by software, that undertakes either individual transformation processes, or links these into processing workflows. A full record of the research undertaken requires preservation of these processing steps and software tools employed, in addition to the datasets upon which they acted.

What is happening now?

Exemplars of repositories for research datasets, software and workflows include Dataverse [21, 10], the Dryad Data Repository [38, 18], and myExperiment, a social network relating to workflows [17, 12].

What are the next steps?

Efforts at archiving, retrieving and citing digital research objects in standardized ways should be closely linked with open data and open-source software publication approaches, and should converge on common standards and practices. Citations to datasets and other digital research objects within publications should be treated on a par with the current treatment of bibliographic citations. Citations to these in the text should be made with a standard reference mark (in-text reference pointer) and the full reference should be given in the reference list of the publication, using a resolvable globally unique identifier (URL, DOI, HDL). Additionally, a formal semantic representation in OWL/RDF of the metadata describing these research objects, their provenance, their relationships to and citations of one another, etc., would be very useful and is now achievable. However, improved tools are required to reduce the labour of creating such metadata.

Openness and What it Implies

3.4 Derive new financially sustainable models of access

The emergence of the open access (OA) publishing model for the traditional scientific product, the journal article, has been a major driver in the emergence of Force11. OA provides the gateway to new modes of scholarly communication, and is the cornerstone that must be promoted and extended if significant change to the scholarly publishing ecosystem is to take place. But OA per se is not enough. It must be shown to be sustainable through new business models, and must be weaved into the academic funding and reward system; neither will be easy. Here is what Force11 advocates to achieve the necessary change to this ecosystem through OA:

- Advocacy for OA through interactions with all the stakeholders mentioned in this document.
- Encouragement of conformance to OA licenses.
- Commitment to make all one's own scholarship as open as possible under the most liberal of those licenses.
- Education of others concerning the features and nuances of OA-based scholarship
- Development of new technologies that assume OA.
- Recognition that OA applies not just to research articles, but also to data, software, bibliographic and citation metadata, books and other components of the scientific process, and the whole scholarly enterprise.
- Recognition that OA applies just as appropriately to emergent research objects.
- Recognition that OA requires sustainable business models, and commitment to work towards achieving those new business models, that are likely to focus less on the content itself and more on the provision of revenue-generating services that facilitate discovery and reuse of that content in ways that advance scholarship.

What is happening now?

The following exemplify that change in the scholarly publishing world is already taking place and is likely to accelerate over time. It is the mandate of Force11 to facilitate that acceleration:

- The increasing number of OA journals, including some that are regarded as comparable with the most highly regarded subscription access publications.
- The emergence of ORCID[3] as a system for creating unique personal identifiers, and hence for author disambiguation and better tagging of all aspects of scholarship.
- The creation of new tools that leverage content e.g. SciVerse[4] and Utopia[5] albeit neither yet in the open access/open source space.
- The development of new article-level metrics and other tools for assessing scholarship.
- The greater sense of awareness to be found within promotion committees concerning the value of alternative forms of scholarship.

What are the next steps?

Force11 members are stakeholders in all aspects of the scholarly enterprise and can influence it in different ways, but all start from the vision outlined above. Some specific steps we now need to take are:

- Start open enterprises that foster change: e.g., new data and software journals, institutional repositories that enable straightforward content exchange.
- Develop tools that highlight non-traditional forms of scholarly output such as database annotations created, blog posts written, and software developed.
- Develop means to assess and highlight the quality of OA content and other non-traditional forms of scholarly output.

3.5 Derive new business models for science publishers and libraries

Current business models for scholarly publication face significant disruption due to many factors: the growth in open access, the advent of alternative publication platforms that exploit new technologies for inexpensive communication and information exchange over the internet, a widening view of what constitutes a publishable research object (e.g. data, workflows), and the challenges of curating, linking and preserving the wider world of digital research objects. Furthermore, it is anticipated that the overall funds dedicated to scholarly communication may well become more restricted in future, at least on a per researcher basis. Both the major customers (research libraries) and brokers (currently, publishers) have an interest in being an active party in shaping the transition to new, sustainable business models, to ensure that the transition is a smooth one.

What is happening now?

The overall market for scholarly communications is on the order of $10 billion per year. The market is not a monolithic one, and disruptions are likely to be somewhat different in different

[3] http://orcid.org/
[4] http://www.hub.sciverse.com/
[5] http://getutopia.com/documents/

disciplines. For example, there is an important distinction between those disciplines where publications are primarily in the form of books rather than journal articles. Also, researchers are growing accustomed to relying on an increasing number of free services. These pose both sustainability risks and opportunities. While freemium services typically manage to recruit only a few percent of users, some of these services can be sustained by a wider marketplace.

Some of these functions face significant challenges. For example, archiving and preservation of research objects, despite its high potential cost, is unique in not directly contributing to reward for producers. For this reason, it will likely be the most difficult to sustainably fund, and may require higher public investment.

What are the next steps?

To be financially viable, new communication modes will need to demonstrate tangible value to both producers and consumers. To be sustainable, the cost recovery streams will need to be aligned to perceived value. An additional factor that should be taken into account is that there are at least three different market sectors to which new products and services may be targeted: tools for producers (aka researchers), enhanced products for consumers (researchers again), and reputation management (for individuals, institutions, and funding bodies).

In Dagstuhl, the Force11 group started to work on a more detailed business model, based on the Business Model Generation methodology [26]. The results of this work will be made available on the Force11 web site [16].

3.6 Derive new methods and metrics for evaluating quality and impact that exploit the technology

Scholarly practices and the way that science is undertaken is changing, as are the possibilities and associated activities of scholarly communication. Yet measures of assessment and impact have not caught up with these changes. Impact is a measure of change. Since these changes can be arbitrarily removed from the immediate outcome, one cannot always easily attribute the changes solely to the action performed. Measuring impact is complex because it depends on context, on purpose, on audience. It can have different effects for different individuals. Similarly, a communication can have different degrees and even polarities of effect. For example, a research paper might be simplified and published by newspapers to make headline news with great societal impact, but be roundly criticized or even ignored by academic colleagues.

What is happening now?

Presently, online versions of 'scholarly outputs' have tended to replicate print forms rather than exploit the affordances and functionalities of the digital terrain. The historical limits of print space are one reason, amongst others, that traditional journal articles tend to represent truncated versions of findings. The assumption is that technology will enable more effective enhanced papers. In addition, scholarly outputs will broaden to include, for example, software tools and social media channels. Work being undertaken under the Alt-metrics[6] umbrella

[6] http://altmetrics.org/manifesto/

pertains here and is to be supported. This has implications for policy. The challenge will then be how to get these metrics accepted by universities, funders and national decision makers.

What are the next steps?

It is accepted that metrics are still needed; however better mechanisms of measurement need to be put in place, that allow for different types of impact and influence.. A multi-dimensional measurement instrument would be useful. It needs to be customisable for specific situations and individual and it must be easy to use both for the individual academic and for the reviewer or decision-maker. What is being measured could include:

- Quality (exploiting new forms of measurement mechanisms).
- Influence (using new forms of alternative metrics).
- Social impact (measured, for example, through development goals).
- Economic impact.
- Contribution to education (use in lectures, reading lists etc.).
- Openness, making scholarly resources shareable, accessible, and re-usable.

Mechanisms for measuring need to be reviewed in an age where traditional forms of peer review are also under critical scrutiny.

Although work has been undertaken to formalise these alternative notions of impact, none are directly applicable today. On the Force11 website, we make some concrete proposals for describing and utilising such new metrics.

4 Related Efforts

The Force11 members have compiled, and will continue to update, a list of others ongoing efforts to improve digital scholarship.[7] You are invited to add to this living document, because we are sure that many other efforts exist, unknown to us. The catalog provides pointers to important papers, relevant blogs government and private sector reports, funding opportunities, policies, domain specific considerations, upcoming and past activities, and organizations.

Relevant papers and books are listed on the Force11 web site [16]. These relate to various aspects of digital scholarship including, but not limited to the reward system, annotation, tools, repositories, text mining, citation of data, textual content in digital form other than research articles (e.g. of eBooks and technical reports), ontologies, metadata standards, semantics, provenance, features of research objects, and workflows.

Additional readings: [1], [2], [3], [4], [6], [7], [8], [9], [15], [19], [24], [25], [28], [30], [32], [33], and [37].

5 Fulfilling this Vision

Force11 has identified the following actions that will contribute towards fulfilling the vision. Some actions apply to all stakeholders, others only to specific groups.

[7] See https://sites.google.com/site/futureofresearchcommunications/links/links.

- Improved collaborative practice, which implies:
  - Increased social media presence.
  - Maximizing informal contacts through conferences, workshops, meetings, calls, webcasts.
  - Joint grant-funded activities leading to the creation of new tools and their description in publications.
  - Other group technology development projects.
- Coordinated standard and technology development, which implies:
  - Wholehearted adoption of W3C web standards and core ontologies.
  - Open source development in response to user specifications from relevant stakeholders.
  - Emphasis on reusability and extensibility.
  - Creation of exemplars which act as drivers for future coordinated efforts, thereby insuring creativity and innovation is part of the development effort; such examples might be:
    - Novel tools that facilitate the use of digital objects.
    - Development of novel metrics to measure non-traditional scholarship.
    - Models for creating useful discipline specific digital repositories.
    - New publishing paradigms.
- Advocacy, which implies:
  - Promoting improved digital scholarship through traditional publication and non-traditional means.
  - Participating in appropriate committees and other organizational bodies that can precipitate change.
  - Fundraising for specific activities in support of change in digital scholarship.

6 Acknowledgements

The Force11 meeting at Schloss Dagstuhl established common ground between interested parties from many different constituencies, and generated a collective resolve to change research communication and e-scholarship for the better. Our mission now is to progress the actions outlined above. We invite you to join the endeavour, and to play a role in shaping the future. There are exciting and rewarding times ahead.

7 Participants

- Bradley P. Allen
Elsevier – Manhattan Beach, US
- Aliaksandr Birukou
CREATE-NET – Povo, Trento, IT
- Judith A. Blake
The Jackson Laboratory – Bar Harbor, US
- Philip E. Bourne
UC San Diego, US
- Simon Buckingham Shum
The Open University – Milton Keynes, GB
- Gully Burns
University of Southern California - Marina del Rey, US
- Leslie Chan
University of Toronto, CA
- Olga Chiarcos
Springer-Verlag – Heidelberg, DE
- Paolo Ciccarese
Harvard University, US
- Timothy W. Clark
Mass General Hospital & Harvard Medical School, US
- Laura Czerniewicz
University of Cape Town, ZA
- Robert Dale
Macquarie University, AU
- Anna De Liddo
The Open University – Milton Keynes, GB
- David De Roure
University of Oxford, GB
- Anita De Waard
Elsevier Labs – Jericho, US
- Stefan Decker
National University of Ireland – Galway, IE
- Alex Garcia Castro
Universität Bremen, DE
- Carole Goble
University of Manchester, GB
- Eve Gray
University of Cape Town, ZA
- Paul Groth
Free University – Amsterdam, NL
- Udo Hahn
Universität Jena, DE
- Ivan Herman
CWI - Amsterdam, NL
- Eduard H. Hovy
Univ. of Southern California – Marina del Rey/ISI, US
- Michael J. Kurtz
Harvard-Smithsonian Center for Astrophysics, US
- Fiona Murphy
Wiley-Blackwell, UK
- Cameron Neylon
Rutherford Appleton Lab. – Didcot, GB
- Steve Pettifer
University of Manchester, GB
- Mike W. Rogers
Europ. Commission Brussels, BE
- David S. H. Rosenthal
Stanford University Libraries, US
- David Shotton
University of Oxford, GB
- Jarkko Siren
Europ. Commission Brussels, BE
- Herbert van de Sompel
Los Alamos National Lab., US
- Peter van den Besselaar
Free Univ. – Amsterdam, NL
- Todd Vision
University of North Carolina – Chapel Hill, US

References

1. Micah Altman and Gary King. A proposed standard for the scholarly citation of quantitative data. *DLib Magazine*, 13(3/4), 2006.
2. R. B. Altman, C. M. Bergman, J. Blake, C. Blaschke, A. Cohen, F. Gannon, and A. Valencia. Text mining for biology—the way forward: opinions from leading scientists. *Genome Biology*, 9 Suppl 2(S7), 2008. doi: 10.1186/gb-2008-9-s2-s7.
3. Teresa K. Attwood, Douglas B. Kell, Philip McDermott, James Marsh, Steve R. Pettifer, and David Thorne. Calling international rescue: knowledge lost in literature and data landslide! *Biochemical Journal*, 424(3):317–333, 2009.
4. Sean Bechhofer, Iain Buchan, David De Roure, Paolo Missier, John Ainsworth, and Carol Goble. Why linked data is not enough for scientists. *Future Generation Computer Systems*, in press.
5. Sean Bechhofer, David De Roure, Matthew Gamble, Carole Goble, and Iain Buchan. Research objects: Towards exchange and reuse of digital knowledge. Paper presented at The Future of the Web for Collaborative Science (FWCS 2010), April 2010, Raleigh, NC, US. http://eprints.ecs.soton.ac.uk/18555/, 2010.
6. Philip E. Bourne. What do i want from the publisher of the future? *PLoS Computational Biology*, 6(5), 2010. e1000787.
7. Jan Brase. Datacite: A global registration agency for research data. Paper presented at COINFO '09: The Fourth International Conference on the Cooperation and Promotion of Information Resources in Science and Technology, 2009.
8. Leslie Chan, Barbara Kirsop, and Subbiah Arunachalam. Towards open and equitable access to research and knowledge for development. *PLoS Med*, 8(3), 2011. e1001016.
9. Paolo Ciccarese, Marco Ocana, Leyla Jael Garcia-Castro, Sudeshna Das, and Tim Clark. An open annotation ontology for science on web 3.0. *BMC Bioinformatics*, 2 Suppl 2(S4), 2011.
10. Merce Crosas. The Dataverse Network: An open-source application for sharing, discovering and preserving data. *D-Lib Magazine*, 17:1–2, 2011. http://www.dlib.org/dlib/january11/crosas/01crosas.html.
11. David De Roure and Carole Goble. Lessons from myexperiment: Research objects for data intensive research. Paper presented at the eScience Workshop 2009, October 15-17, 2009, Pittsburgh, US. http://eprints.ecs.soton.ac.uk/17744/, 2009.
12. David De Roure, Carole Goble, and Robert Stevens. The design and realisation of the myexperiment virtual research environment for social sharing of workflows. *Future Generation Computer Systems*, 25:561–567, 2009. doi: 10.1016/j.future.2008.06.010.
13. Anita de Waard. Hypotheses, evidence and relationships: The hyper approach for representing scientific knowledge claims. Paper presented at the Workshop on Semantic Web Applications in Scientific Discourse (SWASD 2009), co-located with the 8th International Semantic Web Conference (ISWC-2009), Washington DC, USA, 2009.
14. Anita de Waard. From proteins to fairytales: Directions in semantic publishing. *IEEE Intelligent Systems*, 25(2):83–88, 2010. doi: 10.1109/MIS.2010.49.
15. Y. Engestrom. Communication, discourse and activity. *The Communication Review*, 3(1):165–185, 1999.
16. Force11 Community. The Future of Research Communication and e-Scholarship (Force11). http://force11.org/.
17. Carole Anne Goble and David Charles De Roure. myexperiment: social networking for workflow-using e-scientists. Paper presented at the Proceedings of the 2nd workshop on Workflows in support of large-scale science, Monterey, California, USA, 2007.
18. Jane Greenberg. Theoretical considerations of lifecycle modeling: An analysis of the dryad repository demonstrating automatic metadata propagation, inheritance, and value

system adoption. *Cataloging and Classification Quarterly*, 47(3–4):380–402, 2009. doi: 10.1080/01639370902737547.

19 S. A. Greenberg. How citation distortions create unfounded authority: analysis of a citation network. *BMJ*, 339(b2680), 2009. doi: 10.1136/bmj.b2680.

20 Ian Hargreaves. Digital opportunity: A review of intellectual property and growth. Retrieved from http://www.ipo.gov.uk/ipreview-finalreport.pdf, 2011.

21 G. King. An introduction to the dataverse network as an infrastructure for data sharing. *Sociological Methods Research*, 36, May 2007.

22 Leonard Kniffel and Charles W. Bailey. Cuts, freezes widespread in academic libraries. Technical report, American Libraries, 2009. http://www.ala.org/ala/alonline/currentnews/newsarchive/2009/may2009/academiclibrarywoes051309.cfm.

23 Peter Murray-Rust. Data-driven science: a scientist's view, 2007. Paper presented at the NSF/JISC Repositories Workshop, Phoenix AZ, April 10, 2007. http://www.sis.pitt.edu/~repwkshop/papers/murray.html.

24 Cameron Neylon. Open research computation: An ordinary journal with extraordinary aims. Science in the Open, http://cameronneylon.net/blog/open-research-computation-an-ordinary-journal-with-extraordinary-aims/, 2011.

25 Cameron Neylon. Time for total scientific openness. *New Scientist*, 2828, 2011.

26 Alexander Osterwalder and Yves Pigneur. *Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers*. John Wiley and Sons, 2010.

27 David S. H. Rosenthal and Victoria Reich. Archiving supplemental materials. *Information Standards Quarterly*, 22(3), 2010. doi: 10.3789/isqv22n3.2010.04.

28 Robert Sanderson and Herbert Van de Sompel. Open annotation: Beta data model guide. Retrieved 9 September, 2011, from http://www.openannotation.org/spec/beta/, 2011.

29 David M. Schultz. The proliferation of scientific literature. Eloquent Science: http://eloquentscience.com/2011/06/the-proliferation-of-scientific-literature/, 2011.

30 David Shotton. Semantic publishing: the coming revolution in scientific journal publishing. *Learned Publishing*, 22(2):85–94, 2009. doi: 10.1087/2009202.

31 David Shotton. The five stars of online journal articles: an article evaluation framework. *Nature Preceedings*, 17, October 2011.

32 Gavin J. D. Smith, Dhanasekaran Vijaykrishna, Justin Bahl, Samantha J. Lycett, Michael Worobey, Oliver G. Pybus, and Andrew Rambaut. Origins and evolutionary genomics of the 2009 swine-origin h1n1 influenza a epidemic. *Nature*, 459(7250):1122–1125, 2009. doi: 10.1038/nature08182.

33 Susan Leigh Star and James R Griesemer. Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in berkeley's museum of vertebrate zoology, 1907–39. *Social Studies of Science*, 19(3):387–420, 1989.

34 The Economist. A world of hits. *The Economist*, 2009. http://www.economist.com/node/14959982.

35 B. Vickery. A century of scientific and technical information. *Journal of Documentation*, 55:476–527, 1999.

36 http://www.w3.org/wiki/SweoIG/TaskForces/CommunityProjects/LinkingOpenData.

37 Etienne Wenger. Communities of practice and social learning systems. *Organization*, 7(2):225–246, 2000. doi: 10.1177/135050840072002.

38 H. White, S. Carrier, A. Thompson, J. Greenberg, and R. Scherle. The Dryad data repository: A Singapore framework metadata architecture in a DSpace environment. In J. Greenberg and W. Klas, editors, *Proceedings of the International Conference on Dublin Core and Metadata Applications*, pages 157–162, 2008.