DAGSTUHL
REPORTS

Volume 8, Issue 8, August 2018

ISSN 2192-5283

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/dagpub/2192-5283

Publication date
March, 2019

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`.

Digital Object Identifier: 10.4230/DagRep.8.8.i

Aims and Scope
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:

- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e. g. summaries from panel discussions or open problem sessions.

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de
http://www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 18321

Web Application Security

Edited by

Martin Johns[1], Nick Nikiforakis[2], Melanie Volkamer[3], and John Wilander[4]

1 **TU Braunschweig, DE,** mj@martinjohns.com
2 **Stony Brook University, US,** nick@cs.stonybrook.edu
3 **KIT – Karlsruher Institut für Technologie, DE,** melanie.volkamer@kit.edu
4 **Apple Computer Inc. – Cupertino, US,** john@wilander.net

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 18321 "Web Application Security". In this third seminar on the topic, a healthy mix of academics, practitioners and representatives of all major browser vendors reflected on the last decade of web security research and discussed the upcoming security challenges for the Web platform. In addition, for the first time, the list of attendees included several members of the *human factors in security* community, to enable broadening the web security topic towards this important facet of application security.

Seminar August 5–8, 2018 – http://www.dagstuhl.de/18321

2012 ACM Subject Classification Security and privacy → Browser security, Security and privacy → Information flow controlSecurity and privacy → Web protocol security, Security and privacy → Software security engineeringSecurity and privacy → Web application security, Security and privacy → Privacy protections, Security and privacy → Usability in security and privacy

Keywords and phrases Web Application Security, Browser Security, Software Security, Human Aspects in Security

Digital Object Identifier 10.4230/DagRep.8.8.1

1 Executive Summary

Martin Johns
Nick Nikiforakis
Melanie Volkamer
John Wilander

Introduction

Motivation

Since its birth in 1990, the Web has evolved from a simple, stateless delivery mechanism for static hypertext documents to a fully-fledged run-time environment for distributed, multi-party applications. Even today, there is still a continuous demand for new features and capabilities which drives the Web's evolution onwards. This unplanned and often chaotic development has led to several deeply ingrained security and privacy problems that plague the platform:

- The Web's original hypertext, multi-origin nature which is manifested in the design of HTML and HTTP is in fundamental conflict with JavaScript's Same-Origin Policy, the Web's most important security mechanism.
- Important security properties, such as end-to-end communication security or endpoint identity are outside of the control of the actual applications. Instead, they depend on the security of external entities, such as domain name servers or certificate authorities.
- Data/code separation in web applications is practically infeasible, as the HTTP link between server-side application logic and client-side application interface requires an intermixing of protocol, data and code fragments within a single continuous character stream.
- HTTP is a stateless protocol without a native session or authentication tracking concept.
- Users are not aware of general or application specific threats. Protecting against these threats (incl. to know which security indicators to trust) is nowadays difficult and time consuming.

Using this fragile basis, critical applications are created, that long have left the strict client-server paradigm, on which the Web was initially built. Instead, scenarios are realized that involve several mutually distrusting entities in a single security and application context. In many cases the browser is the link that connects the remote parties, either via direct JavaScript inclusion, web mashups, or through the usage of web protocols, such as OpenID and OAuth.

The accumulated ballast of the last two decades of web evolution, the ever growing functional demands of sophisticated web applications and the ambitious vision of the web platform's drivers creates an exciting tension field which is in constant conflict with the required security assurances of high value business applications.

Since approximately ten years, academic security and privacy research has recognized the importance of the web platform and the unique characteristics and challenges of the web security and privacy topic. And while specific techniques, that originated from academic research, such as the Content Security Policy, have been adapted in practice, the fundamental security problems of the web remain and the overall vulnerability landscape is getting worse, as it can be seen in the constant flow of reported web security issues in bug trackers and vulnerability databases.

Academic web security research has started 2007 and usable security research started almost at the same time. In the context of this Dagstuhl Seminar, we will revisit the lessons learned from the last decade and revisit the success stories and mistakes that have been made. Questions, that have to be raised in include "What has worked?", "What has been taken up by industry?", "What failed and why?", and – most importantly – "What did we learn?"

Seminar Objectives

Today, several unconnected groups drive the topic, including Security, Privacy as well as Usable Security & Privacy Academics, standardization, and browser vendors. The seminar will facilitate essential exchange between them. This will allow academia to directly influence browser vendors and standardization representatives, and allow industry representatives to influence the research community.

Overview

Participants

The seminar was well attended with 39 participants. A good balance of European and American researchers was present. Furthermore, the group represented a nice mix of participants of academia and industry. Compared to the previous editions, not only researchers from the web security area participated but also from the field of human factors in security.

Structure

This was the third Dagstuhl seminar on Web application security. The seminar's organisation combined overview presentation of various subfields, highlight talks, and discussions in working groups. In particular the overview presentations were important to connect the two research fields web security from a more technical point of view and human factors in security. This way, also a good, comprehensive view on current activities and open problems in the realm of Web application security in particular from a user's point of view could be achieved and areas for potential future collaborations could be identified.

Summary

Talks

The following people presented either an overview of their research field, very recent research results or overarching observations on the field of web application security. Please also refer to Section 3 for selected talk abstracts.

- Stefano Calzavara, University of Venezia, IT: REASON – A programmable architecture for secure browsing
- Luca Compagna, SAP Labs France – Mougins, FR: Analysis & Detection of Authentication Cross-Site Request Forgeries
- Lieven Desmet, KU Leuven, BE: Detecting and Preventing Malicious Domain Registrations in the .eu TLD
- Steven Englehardt, Mozilla – Mountain View, US: No Boundaries: Data exfiltration by directly embedded tracking scripts
- Thomas Gross, Newcastle University, GB: Investigating Cognitive and Affective Predictors Impacting Password Choice
- Mario Heiderich, Cure53 – Berlin, DE, DOMPurify: Client-Side Protection Against XSS and Markup Injection
- Boris Hemkemeier, Commerzbank AG – Frankfurt, DE: Web application security in vulnerable environments
- Martin Johns, TU Braunschweig, DE: WebAppSec @ Dagstuhl – The Third Iteration
- Christoph Kerschbaumer, Mozilla – San Francisco, US: Could we use Information Flow Tracking to generate more sophisticated blacklists?
- Pierre Laperdrix, Stony Brook University, US: Browser fingerprinting: current state and possible future
- Sebastian Lekies, Google Switzerland – Zürich, CH: Trusted Types: Prevent XSS with this one simple trick!
- Benjamin Livshits, Imperial College London, GB: Browser Extensions for the Web of Value

- Marius Musch, TU Braunschweig, DE: On measurement studies and reproducibility
- Lukasz Olejnik, Independent researcher, W3C TAG, FR: Private browsing modes guaranteed. On the example of Payment Request API
- Juan David Parra, Universität Passau, DE: Computational Resource Abuse through the Browser
- Giancarlo Pellegrino, Stanford University, US: Removing Browsers from the Equation: A New Direction for Web Application Security
- Tamara Rezk, INRIA Sophia Antipolis, FR: Content Security Policy Challenges
- Konrad Rieck, TU Braunschweig, DE: Beyond the Hype: Web Security and Machine Learning?
- Andrei Sabelfeld, Chalmers University of Technology – Göteborg, SE: A Challenge for Web of Things: Securing IoT Apps
- Sebastian Schinzel, FH Münster, DE: Handling HTML Emails after the Efail Attacks
- Zubair Shafiq, University of Iowa – Iowa City, US: The Arms Race between Ad Tech vs. Adblockers: Key Challenges and Opportunities
- Lynsay Shepherd, Abertay University – Dundee, GB: How to Design Browser Security and Privacy Alerts
- Dolière Francis Somé, INRIA Sophia Antipolis, FR: The Same Origin Policy and Browser Extensions
- Ben Stock, CISPA – Saarbrücken, DE: Persistent Client-Side Cross-Site Scripting in the Wild
- Melanie Volkamer, KIT – Karlsruher Institut für Technologie, DE: Web Security Meets Human Factors in Security
- Mike West, Google – München, DE: HTTP State Tokens

Conclusions

This seminar was the third Dagstuhl Seminar von Web Application Security, following Seminar 09141 (2009) and Seminar 12401 (2012). Thus, it was a great opportunity to reflect on a decade of web security research. In 2009 the field was largely undefined and that year's seminar offered a wild mix of various topics, some with lasting impact and many that went nowhere. Where the 2009 seminar was overly broad, the 2012 iteration had a comparatively narrow focus as the seminar was dominated by the notion that solving web security mainly revolves around solving the security properties of JavaScript.

This year's seminar reflected the ongoing maturing of the topic very well. Fundamental problems, such as Cross-site Scripting or the Web Browser security model, are well explored and their understanding served as a great foundation for the seminar's discussions. This allowed the extension of the topic toward important facets, such as privacy problems or human factors. While the addressed topics were too broad and the time for overarching discussions was limited due to the three-day format of the seminar, the sparked discussions were fruitful for several follow-up activities (see above). An underlying theme of the seminar can be summarized as "the last decade of web security has broad good progress and development but the overall problem is still neither fully understood nor solved". Especially, the newly introduced dimension of integrating human factors in security, which was reflected through including several high-profile members of this community in the seminar, is still immature.

One of the seminar's prime objectives has been reached very nicely: The fostering of collaboration between the different web security communities. For one, several compelling

interactions between practitioners from industry (such as SAP, Commerzbank and Cure53) and researcher from academia took place. Furthermore, thanks to the fact that all major web browser vendors (plus the new privacy-centric browser Brave) were represented at the seminar, both cross-browser vendor interaction as well as browser/academia collaborations were initiated, with the browser-based sanitizer initiative (see breakout session 4.3) being a prominent example.

2 Table of Contents

3 Overview of Talks

3.1 No Boundaries: Measuring data exfiltration by third-party scripts

Steven Englehardt (Mozilla – Mountain View, US)

Web tracking is pervasive. A core requirement of web tracking – the identification of individuals across website – is increasingly difficult as browser vendors adopt strict cookie policies and users take steps to protect their privacy. As a result, web trackers have deployed invasive tracking techniques that lack user (and sometimes browser) controls.

In this talk I'll explore findings from our recent web tracking measurements, which show the lengths to which trackers have gone to collect user information. Examples include: the abuse of browser autofill to collect email addresses, the exfiltration of information from social login APIs, and the collection of user information from the DOM. We find that some websites which embed these trackers are – much like users – completely unaware of these practices. I'll close with a discussion of our options for preventing this type of tracking.

3.2 Could we use an Information Flow Tracking to generate more sophisticated blacklists?

Christoph Kerschbaumer (Mozilla – San Francisco, US)

JavaScript (JS) has become the dominant programming language of the Internet and powers virtually every web page. User agents face a difficult situation: on the one hand JavaScript allows websites to provide a rich user experience; on the other hand JavaScript allows adversaries to perform malicious actions. To distinguish between good and malicious JavaScript at runtime has proven complicated and quite often browser vendors see no other options than relying on pre-rendered blacklists to block malicious JavaScript from executing.

While the approach of building an Information Flow Tracking system into a web browser has proven questionable: (a) because of the performance drawback, and (b) because of various loopholes which do not allow precise information tracking in a browser mostly due to JavaScripts dynamic nature. Nevertheless, an enhanced browser performing information flow tracking might still be able to detect malicious actions of JavaScript and hence provide input for creating more sophisticated blacklists.

Hence we ask: Could we use an Information Flow Tracking to generate more sophisticated blacklists?

3.3 Browser fingerprinting: current state and possible future

Pierre Laperdrix (Stony Brook University, US)

Joint work of Vastel, Antoine; Rudametkin, Walter; Rouvoy, Romain; Gómez-Boix Alejandro; Baudry, Benoit
Main reference Alejandro Gómez-Boix, Pierre Laperdrix, Benoit Baudry: "Hiding in the Crowd: an Analysis of the Effectiveness of Browser Fingerprinting at Large Scale", in Proc. of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018, pp. 309–318, ACM, 2018.
URL https://doi.org/10.1145/3178876.3186097

After a brief introduction on what browser fingerprinting is, we will take a look at the latest studies published in the domain* and the current ecosystem regarding fingerprinting protection. Then, we will see what lies ahead by talking about how this technique could be used positively to increase online security.

Open questions: Is there a future for constructive fingerprinting? If so, how?

*3 papers:

- FP-Scanner: The Privacy Implications of Browser Fingerprint Inconsistencies Antoine Vastel, Pierre Laperdrix, Walter Rudametkin, Romain Rouvoy (USENIX Sec. 2018)
- FP-STALKER: Tracking Browser Fingerprint Evolutions Antoine Vastel, Pierre Laperdrix, Walter Rudametkin, Romain Rouvoy (S&P 2018)
- Hiding in the Crowd: an Analysis of the Effectiveness of Browser Fingerprinting at Large Scale Alejandro Gómez-Boix, Pierre Laperdrix, Benoit Baudry (WWW 2018)

3.4 Security of Modern Mobile Browsers

Nick Nikiforakis (Stony Brook University, US)

Joint work of Meng Luo, Oleksii Starov, Nima Honarmand, Nick Nikiforakis
Main reference Meng Luo, Oleksii Starov, Nima Honarmand, Nick Nikiforakis: "Hindsight: Understanding the Evolution of UI Vulnerabilities in Mobile Browsers", in Proc. of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017, pp. 149–162, ACM, 2017.
URL https://doi.org/10.1145/3133956.3133987

Much of recent research on mobile security has focused on malicious applications. Although mobile devices have powerful browsers that are commonly used by users and are vulnerable to at least as many attacks as their desktop counterparts, mobile web security has not received the attention that it deserves from the community. In particular, there is no longitudinal study that investigates the evolution of mobile browser vulnerabilities over the diverse set of browsers that are available out there. In this paper, we undertake the first such study, focusing on UI vulnerabilities among mobile browsers. We investigate and quantify vulnerabilities to 27 UI-related attacks–compiled from previous work and augmented with new variations of our own–across 128 browser families and 2,324 individual browser versions spanning a period of more than 5 years. In the process, we collect an extensive dataset of browser versions, old and new, from multiple sources. We also design and implement a browser-agnostic testing framework, called Hindsight, to automatically expose browsers to attacks and evaluate their vulnerabilities. We use Hindsight to conduct the tens of thousands of individual attacks that were needed for this study. We discover that 98.6% of the tested browsers are vulnerable to at least one of our attacks and that the average mobile web browser is becoming less secure with each passing year. Overall, our findings support the conclusion that mobile web security has been ignored by the community and must receive more attention.

3.5 Beyond the Hype: Web Security and Machine Learning?

Konrad Rieck (TU Braunschweig, DE)

Machine learning has made considerable progress in the last years. Unfortunately, this progress is overshadowed by a hype in the industry, and it has become difficult to separate good ideas from marketing phrases. While this talk cannot solve this problem, it aims at highlighting three recent learning concepts that might be fruitful in the context of Web security and deserve to be discussed, irrespective of the current hype.

3.6 A Challenge for Web of Things: Securing IoT Apps

Andrei Sabelfeld (Chalmers University of Technology – Göteborg, SE)

Joint work of Iulia Bastys, Musard Balliu, Andrei Sabelfeld
Main reference Iulia Bastys, Musard Balliu, Andrei Sabelfeld: "If This Then What?: Controlling Flows in IoT Apps", in Proc. of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018, pp. 1102–1119, ACM, 2018.
URL https://doi.org/10.1145/3243734.3243841

IoT apps empower users by connecting a variety of otherwise unconnected services. Unfortunately, the power of IoT apps can be abused by malicious makers, unnoticeably to users. We demonstrate that popular web-based IoT app platforms are susceptible to several classes of attacks that violate user privacy, integrity, and availability. We estimate the impact of these attacks by an empirical study. We suggest short/medium-term countermeasures based on fine-grained access control and long-term countermeasures based on information flow tracking. Finally, we discuss general trends and challenges for securing the Web of Things.

3.7 Efail: Breaking S/MIME and OpenPGP Email Encryption using Exfiltration Channels

Sebastian Schinzel (FH Münster, DE)

Joint work of Damian Poddebniak, Christian Dresen, Jens Müller, Fabian Ising, Sebastian Schinzel, Simon Friedberger, Juraj Somorovsky, Jörg Schwenk
Main reference Damian Poddebniak, Christian Dresen, Jens Müller, Fabian Ising, Sebastian Schinzel, Simon Friedberger, Juraj Somorovsky, Jörg Schwenk: "Efail: Breaking S/MIME and OpenPGP Email Encryption using Exfiltration Channels", in Proc. of the 27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, August 15-17, 2018., pp. 549–566, USENIX Association, 2018.
URL https://www.usenix.org/conference/usenixsecurity18/presentation/poddebniak

The Efail attack abuses malleable encryption in the respective modes of encryption in the OpenPGP and S/MIME standards. The attacker changes an existing ciphertext in a way that its plaintext is exfiltrated to the attacker when opened. For encrypted emails, the attacker edges the actual content of the email in HTML tags that perform external HTTP requests (backchannels). The victim's email client decrypts the email and sends the plaintext to the attacker. Outdated cryptography clearly is the culprit here and deploying authenticated

encryption (AE) ciphers to the standards could prevent this attack in the future. Besides this cryptographic weakness, HTML emails also play an important role.

While it is possible to port Efail-like attacks to any data standard supporting backchannels, HTML makes the attack particularly easy. HTML emails and especially remote content loading (e.g. images, style sheets) can be used for user tracking and were known to be a privacy issue for many years. While it is quite common for privacy advocates to disable HTML in emails completely, most non-technical users insist on HTML emails because they value rich typesetting in their day-to-day work. This raises some questions:

Is HTML the way to go for future typesetting of emails? Are there safer alternatives?

What is a safe subset of the HTML standards that allows rich typesetting, but without allowing user-tracking or Efail-like attacks?

How to enforce this safe subset in existing emails clients?

3.8 How to Design Browser Security and Privacy Alerts

Lynsay Shepherd (Abertay University – Dundee, GB) and Karen Renaud (University of Abertay – Dundee, GB)

Main reference Lynsay A. Shepherd, Karen Renaud: "How to design browser security and privacy alerts", CoRR, Vol. abs/1806.05426, 2018.
URL http://arxiv.org/abs/1806.05426

Browser security and privacy alerts must be designed to ensure they are of value to the end-user, and communicate risks efficiently. We performed a systematic literature review, producing a list of guidelines from the research. Papers were analysed quantitatively and qualitatively to formulate a comprehensive set of guidelines. Our findings seek to provide developers and designers with guidance as to how to construct security and privacy alerts. We conclude by providing an alert template, highlighting its adherence to the derived guidelines.

3.9 REASON – A programmable architecture for secure browsing

Stefano Calzavara

Joint work of Stefano Calzavara, Riccardo Focardi, Niklas Grimm, Matteo Maffei
Main reference Stefano Calzavara, Riccardo Focardi, Niklas Grimm, Matteo Maffei: "Micro-policies for Web Session Security", in Proc. of the IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016, pp. 179–193, IEEE Computer Society, 2016.
URL https://doi.org/10.1109/CSF.2016.20

The REASON project is a research proposal which I wrote with the goal of improving the security architecture of web browsers. More specifically, REASON aims at replacing the traditional Same Origin Policy (SOP) of web browsers with a programmable security monitor amenable for formal verification.

Preliminary evidence of the effectiveness of the proposal was given in a paper at CSF'16, where a small fragment of the architecture was designed and implemented. This talk will discuss the main motivations behind REASON, its benefits and a few ideas on how to implement it on top of existing browsers.

3.10 Persistent Client-Side Cross-Site Scripting in the Wild

Ben Stock (CISPA – Saarbrücken, DE)

Joint work of Ben Stock, Marius Steffens

The Web has become highly interactive and an important driver for modern life, enabling information retrieval, social exchange, and online shopping. From the security perspective, Cross-Site Scripting (XSS) is one of the most nefarious attacks against Web clients. XSS was long since believed to fall into three categories: reflected, persistent, or DOM-based XSS. In this paper, we present the first systematic study of the threat of Persistent Client-Side XSS, which lies in the intersection of persistent and DOM-based XSS. While the existence of this class has been acknowledged, especially by the non-academic community like OWASP, prior works have either only found such flaws as side effects of other analyses or focussed on a limited set of applications to analyze. Therefore, the community lacks in-depth knowledge about the actual prevalence of Persistent Client-Side XSS.

To close this research gap, we leverage taint tracking to identify suspicious flows from client-side persistent storage (Web Storage, cookies) to dangerous sinks (HTML, JavaScript, etc.). We discuss two attacker models capable of injecting malicious payloads into these storages: one that can manipulate HTTP communication (e.g., in a public WiFi), another that abuses existing reflected Client-Side XSS vulnerabilities to persist their payload. With our tainting methodology and these models in mind, we study the prevalence of Persistent Client-Side XSS in the Alexa Top 5,000 domains. We find that more than 8% of them have unfiltered data flows from persistence to a dangerous sink, which showcases the developers' inherent trust in the integrity of storage content. Investigating those vulnerable flows allows us to categorize them into four disjoint categories and propose appropriate mitigations.

3.11 Human Factors in Web Application Security (and Privacy)

Melanie Volkamer (KIT – Karlsruher Institut für Technologie, DE)

Joint work of currenty and previous members of the SECUSO research group as well as Karen Renaud

The talk starts of by explaining main goals in the research area of human factors in security and privacy as well as the main ideas behind the human centered security / privacy by design methodology including the importance of identifying users' mental models and acknowledging that security / privacy is usually not the users primary task. Then selected research results in the area of web application security are presented: This includes just in time and place security interventions to support users in avoiding to provide sensitive information on http pages [1] and to support them in checking links in emails before actually clicking the link [2]. It also includes proposals how to design UIs for security and privacy settings [3]. The talk concludes by raising open research questions in this area.

References

1 Melanie Volkamer, Karen Renaud, Gamze Canova, Benjamin Reinheimer, Kristoffer Braun. *Design and Field Evaluation of PassSec: Raising and Sustaining Web Surfer Risk Awareness*. TRUST 2015: 104-122, Springer, 2015

2 Melanie Volkamer, Karen Renaud, Benjamin Reinheimer, Alexandra Kunz. *User experiences of TORPEDO: TOoltip-poweRed Phishing Email DetectiOn*. Computers and Security 71: 100-113, 2017

3 Oksana Kulyk, Peter Mayer, Oliver Käfer, Melanie Volkamer. *A Concept and Evaluation of Usable and Fine-Grained Privacy-Friendly Cookie Settings Interface*. IEEE TrustCom 2018

4 Break Out Sessions

We had two time slots to discuss in small groups and two to present and discuss the results of these break out sessions with the entire group. The following topics were discussed.

- Browser Fingerprinting: Friend or Foe?
- Policies and Capabilities
- Cookies are Bad (for Authentication)
- My Browser Needs a Sanitizer
- Browser Warning Fatigue
- Browser Extensions

In the following sections, we will briefly document the individual sessions' discussions and results.

Browser Fingerprinting: Friend or Foe?

This breakout session covered various topics under the umbrella of browser fingerprinting.

Destructive vs. Constructive Use. Traditionally, browser fingerprinting has been treated by researchers and privacy-aware users as an intrusive practice that should, ideally, be detected and stopped. Yet the act of detecting and recognizing the device of a user can be used for constructive purposes, i.e., detecting the takeover of an account by the fact that the current user's fingerprint does not match the fingerprint collect during previous visits.

The participants of this breakout session discussed the advantages and disadvantages of using browser fingerprinting as an intrusion detection technique. Some participants mentioned that there already exist companies that provide bot-detection services based on device fingerprinting by attempting to recognize a device fingerprint as belonging to a popular bot/attack tool. Others argued that this would be a losing strategy in the long run, since attackers could randomize the fingerprint of their bot so that it stops matching the previously recorded fingerprints. There was some level of disagreement in terms of how feasible this is in the long run, as defenders only need to find one feature to recognize the true nature of a bot.

During this discussion of constructive (improving security) vs. destructive (worsening privacy) fingerprinting, some participants mentioned that if fingerprinting is used constructively, perhaps it can be limited to first-party websites, i.e., the fingerprinting script should be collected and used by the website that a user visits, and not by third parties present on that website.

Diversity and linkage of fingerprints. Most past studies involving browser fingerprinting with the help of volunteers, have reported a high-level of fingerprinting uniqueness, i.e., different users exhibiting different fingerprints which can be used to tell them apart. One of the participants mentioned that in a recent paper published in WWW 2018, researchers used a popular website to deploy their fingerprinting code and were thus able to collect fingerprints from the "general public" rather than those of privacy-aware volunteers. This study discovered that only 30% of users was unique, which is significantly less than prior studies which reported 90% or more uniqueness. The participants concluded that we need more research to identify the main reasons behind this lack of uniqueness reported by this recent stud.

User control. Given the group's discussion of the constructive use of browser fingerprinting, some participants highlighted that current fingerprinting is done in a surreptitious manner which makes people further distrust it. That is, JavaScript programs collect user information and create fingerprints of the user's browsing environment without the knowledge or consent of users.

Some participants, proposed bringing fingerprinting "to the surface" by asking users whether they want to be fingerprinted (similar to current browser popups related to geolocation and web notifications). By making this choice explicit, these participants argued that users could learn to trust a certain number of websites with their fingerprints (by accepting the relevant dialogues) allowing browser vendors and researchers to defend against browser fingerprinting that is done surreptitiously and without user consent.

4.1 Policies and Capabilities

This breakout session covered security and privacy policies set for pages and contexts, and the capabilities of JavaScript in a specific context.

Policies. Content Security Policy (CSP) has an interesting scope in that it limits code injection but not markup injection. Should we expand on CSP or come up with a complimentary policy mechanism? Further, a threat model is missing from CSP. It is mostly about avoiding injection and doesn't address data exfiltration. Was this intentional? Exfiltration can happen in many ways that are not URL-based resource loads such as window.open() + postMessage() and window.name. CSP is also used as mixed content protection and to avoid third-party script inclusion by your own developers. Going the other direction, should we create a strict CSP that is a fragment of CSP for specifically fighting XSS? We could have similar fragments for controlling framing.

4.1.1 Capabilities

Today, all JavaScript in an execution context are created equal. The origin of them or whether they are inline or file-based doesn't affect their powers over content, state, and network traffic. Could we restriction JavaScript use of password fields, payment APIs, computational resources, fingerprinting vectors etc to only a trusted subset?

If we have frame separation (cross-origin or not) we could support a CPU policy per frame. We could invent a new restricted script tag, for instance for ad scripts. Responsible (or previously compromised) sites will use this for third parties. These scripts would then be restricted in the ways described above. Or could some JavaScript sandbox be what we want?

A problem here is so called script gadgets which are very prevalent. They allow for ROP-style malicious code injection by inserting specific elements into the DOM that trigger code paths in legitimate libraries/frameworks (with full powers). This can be leveraged to cross the restriction boundary. Iframes and the sandbox directive may be too restrictive today. It's scripting on or off.

With a new script element rather than attribute on current script elements would allow us to get out of the gadget mess since vulnerable libraries/frameworks will not have flaws for the new script element. An alternative would be to ship something like <script capabilities="ad"></script>, wait a year later, then require it or block based on a blacklist of ad tech origins.

4.2 Cookies are Bad (for Authentication)

Cookies are primary targets of security and privacy attacks such as cross-site scripting, rogue scripting, and speculative execution attacks such as Spectre. This breakout session aimed at looking at how cookies are used today and seeing if we can achieve the same functionality with something more secure.

The current state of cookie usage. Recent statistics of cookie usage in Google Chrome:

- HttpOnly cookies $\approx 9\%$
- Secure cookies $\approx 7\%$
- SameSite cookies $\approx 0.03\%$

This shows how low the adoption of security measures are for this important protocol feature. In addition, websites use up to 180 cookies per site and up to 4kb per cookie which hurts network performance significantly.

Cookie purposes today:

- Hold authentication state.
- User recall (know that a series of requests are from the same user agent).
- Ad (re)targeting.
- Ad/click attribution.
- On-device storage.
- User preference (UI choices or other web app settings).

Towards a better authentication mechanism. We would like to deprecate cookies for the purpose of authentication/user identification in browsers, not for HTTP in general. To get there, these two things were mentioned:

1. Drive down the use of plaintext cookies is good.
2. Drive down the JavaScript use of cookies is good.

The rest of the session focused on Mike West's proposal for a different protocol state mechanism: https://mikewest.github.io/http-state-tokens/ which was discussed in depth.

Migration to a new mechanism. If we were to move to such a mechanism, how would be deprecate cookies, at least for authentication purposes in web browsers?

1. Introduce it.
2. Encourage usage.
3. Now we've given developers an alternative and can start removing cookie support.

A final note on migration was that maybe the browser should not send a token on the first page load, but instead have the server to opt in. The browser could announce its support for the token mechanism.

4.3 My Browser Needs a Sanitizer

The session was joined by participants from Google, Microsoft, Mozilla, SAP and Cure53. The goal was to evaluate whether a browser should expose a HTML Sanitizer API or if this should rather be done by external JavaScript libraries.

The participants agreed that indeed the browser should indeed be the one to offer that feature for a variety of reasons. The discussion then focused on challenges and possible limitations and, as a result, the participants agreed on the next steps.

Those steps are as follows:

1. Creation of a proposal for WICG – essentially the authoring of an "explainer doc"
2. The initiation of authoring a specification draft and further discussions.

The "explainer doc" has by now been published, the spec is in preparation and is being authored by Mozilla and Cure53.

4.4 Browser Warning Fatigue

The group first worked on a common understanding of different types of browser "warnings" with different characteristics, e.g. (1) there are those which force you to make a decision (blocking) and those that appear more like a notice (you don't have to make a decision); (2) there are those that appear as icon (e.g., the lock icon; you may get more information when clicking on the icon) and those that contain text (e.g. explaining the situation why this warning now is shown and what the user needs to decide on); (3) there are those provided by the browser and those from the visited webpage; correspondingly also the positioning varies. For the webpage one the question whether tick boxes next to statements that one agrees on (privacy) policies should be considered as 'warning' was discussed.

With respect to the issues with various types of browser warnings, participants discussed the often mentioned habituation issue. The question was whether this issue is a consequence of badly designed warnings appearing too often without any consequence when ignoring them or whether habituation is an issue of any warning and better design will not help. It was agreed that it should be possible to not just decide this one situation but to tell the system that similar situations should be decided automatically the same way without being actively interrupted again in future. It was discussed whether it is possible to predict user's decisions on warning dialogues (in particular in the privacy context) and therefore make the decisions automatically (or at least provide an option for the user that these decisions can be make automatically).

It was furthermore agreed on that warnings in terms of asking the user to decide should only be displayed if users can make an informed decision based on the information provided in the warning.

There was also a discussion on evaluating warnings in particular wrt to whether they cause fatigue. The issue with fatigue is that you may only measure it after people having used a system with the to be evaluated warnings for some time; which means one need to go for a field study but making sure that the underlying system does not introduce any security issues for the participants.

4.5 Browser Extensions

The final break-out session was dedicated to the topic of browser extensions. In this context several orthogonal topics were touched upon.

Permission System. Similar to mobile apps, browsers utilize permission systems to mitigate potential security problems by malicious extensions. Unfortunately, due to the technical intrinsic of the web model, the current permission granularity is insufficient. For instance, in many cases, such as DOM or Network access, the technically available options are to coarse, being essential full or no access. Within the sessions, alternative approaches were discussed, including moving away from tying permissions to technical capabilities and instead moving to activities.

Extension Vetting. A joint cross-vendor approach toward unified vetting of extension was proposed, as – thanks to standards such as the web extension model – an increasing number of extensions are written that simultaneous target multiple browsers.

Protection Users against malicious extensions. Finally, the session addressed methods to support users (and sites) against malicious extensions. In this context, the notion of trust-classes for web sites was brought up. This would allow the disabling of extensions for security sensitive sites, such as banks, will enabling them on sites with lesser security requirements, such as entertainment sites.

4.6 Aftermath

The seminar was perceived as highly inspiring by the participants. In consequence, it had a fertilizing effect on follow-up activities: Besides various informal collaborations that resulted from discussions in Dagstuhl, we would like to single out results which directly can be attributed to the seminar:

- Upcoming paper on hybrid static/dynamic security analysis of web applications
- Various co-supervised students
- Several research visits (e.g., KIT/Abertay University)
- Several ongoing academic-industry collaborations (e.g., SAP/TU Braunschweig)
- Initiation of a cross-browser specification on a web browser-based API for security handling of untrusted data.

Participants

- Frederik Braun
Mozilla – Berlin, DE
- Achim D. Brucker
University of Sheffield, GB
- Stefano Calzavara
University of Venezia, IT
- Luca Compagna
SAP Labs France – Mougins, FR
- Lieven Desmet
KU Leuven, BE
- Steven Englehardt
Mozilla – Mountain View, US
- Thomas Gross
Newcastle University, GB
- Marian Harbach
Audi AG – Ingolstadt, DE
- Daniel Hausknecht
Chalmers University of Technology – Göteborg, SE
- John Hazen
Microsoft Corporation – Redmond, US
- Mario Heiderich
Cure53 – Berlin, DE
- Boris Hemkemeier
Commerzbank AG – Frankfurt, DE
- Martin Johns
TU Braunschweig, DE
- Christoph Kerschbaumer
Mozilla – San Francisco, US
- Pierre Laperdrix
Stony Brook University, US
- Sebastian Lekies
Google Switzerland – Zürich, CH
- Benjamin Livshits
Imperial College London, GB
- Matteo Maffei
TU Wien, AT
- Marius Musch
TU Braunschweig, DE
- Nick Nikiforakis
Stony Brook University, US
- Lukasz Olejnik
Independent researcher, W3C TAG, FR
- Juan David Parra
Universität Passau, DE
- Giancarlo Pellegrino
Stanford University, US
- Karen Renaud
University of Abertay – Dundee, GB
- Tamara Rezk
INRIA Sophia Antipolis, FR
- Konrad Rieck
TU Braunschweig, DE
- Andrei Sabelfeld
Chalmers University of Technology – Göteborg, SE
- Sebastian Schinzel
FH Münster, DE
- Zubair Shafiq
University of Iowa – Iowa City, US
- Lynsay Shepherd
Abertay University – Dundee, GB
- Dolière Francis Somé
INRIA Sophia Antipolis, FR
- Ben Stock
CISPA – Saarbrücken, DE
- Daniel Veditz
Mozilla – Mountain View, US
- Melanie Volkamer
KIT – Karlsruher Institut für Technologie, DE
- Malte Wedel
SAP SE – Walldorf, DE
- Rigo Wenning
W3C / ERCIM, FR
- Mike West
Google – München, DE
- John Wilander
Apple Computer Inc. – Cupertino, US
- Henrik Willert
1&1 Internet SE – Karlsruhe, DE

Report from Dagstuhl Seminar 18322

Human-Computer Integration

Edited by

Florian Mueller[1], Pattie Maes[2], and Jonathan Grudin[3]

1 **RMIT University – Melbourne, AU,** floyd@floydmueller.com
2 **MIT – Cambridge, US,** pattie@media.mit.edu
3 **Microsoft Research – Redmond, US,** jgrudin@microsoft.com

Abstract

The rise of technology that supports a partnership between user and computer highlights an opportunity for a new era of "human-computer integration", contrasting the previously dominant paradigm of computers functioning as tools. However, most work around these technologies only focused on the instrumental perspective to achieve extrinsic performance objectives. However, phenomenology emphasizes that it is also important to support the experiential perspective, which indicates that technology should also help people pay attention to their lived experiences and personal growth in order to deepen their understanding of their own bodies.

This seminar focuses on embodied integration, where a computer tightly integrates with the person's body. Although an increasing number of systems are emerging, a thorough understanding of how to design such systems is notably absent. The reason for this is the limited knowledge about how such embodied partnerships unfold, and what underlying theory could guide such developments. This seminar brought together leading experts from industry and academia, including those who are central to the development of products and ideas such as wearables, on-body robotics, and exertion systems. The goal was to address key questions around the design of embodied integration and to jump-start collaborations to pioneer new approaches for a human-computer integrated future.

Seminar August 5–10, 2018 – http://www.dagstuhl.de/18322
2012 ACM Subject Classification Human-centered computing → Interaction paradigms
Keywords and phrases Human-computer integration, whole-body interaction, ubiquitous computing, wearables
Digital Object Identifier 10.4230/DagRep.8.8.18
Edited in cooperation with Zhuying Li

1 Executive Summary

Florian Mueller (RMIT University, Melbourne, floyd@exertiongameslab.org)
Jonathan Grudin (Microsoft Research, Redmond, US, jgrudin@microsoft.com)
Pattie Maes (MIT, Cambridge, US, pattie@media.mit.edu)
Zhuying Li (RMIT University, Melbourne, zhuying@exertiongameslab.org)

The rise of technology that supports a partnership between user and computer highlights an opportunity for a new era of "human-computer integration", contrasting the previously dominant paradigm of computers functioning as tools. However, most work around these technologies only focused on the instrumental perspective to achieve extrinsic performance objectives. However, phenomenology emphasizes that it is also important to support the

experiential perspective, which indicates that technology should also help people pay attention to their lived experiences and personal growth in order to deepen their understanding of their own bodies. This seminar focuses on embodied integration, where a computer tightly integrates with the person's body. Although an increasing number of systems are emerging, a thorough understanding of how to design such systems is notably absent. The reason for this is the limited knowledge about how such embodied partnerships unfold, and what underlying theory could guide such developments. This seminar brought together leading experts from industry and academia, including those who are central to the development of products and ideas such as wearables, on-body robotics, and exertion systems. The goal was to address key questions around the design of embodied integration and to jump-start collaborations to pioneer new approaches for a human-computer integrated future.

2 Table of Contents

3 Introduction

Florian Mueller (RMIT University, Melbourne, floyd@exertiongameslab.org)
Jonathan Grudin (Microsoft Research, Redmond, US, jgrudin@microsoft.com)
Pattie Maes (MIT, Cambridge, US, pattie@media.mit.edu)
Zhuying Li (RMIT University, Melbourne, zhuying@exertiongameslab.org)

In Aug 2018, 28 researchers and academics from Europe, North America, and Asia Pacific gathered for a week to discuss the future of Human-Computer Integration (HInt). The goal of the seminar was to discuss the future of what it means to design interactive systems that integrate the human body and technology, a trend highlighted by emerging technologies such as implantables and ingestibles that blurs the boundary of computers and users. The motivation for the seminar stemmed from the realization that until today, with the accelerating technological capabilities, an increasing number of real-world deployments, and growing realizations of ethical and societal implications, it is increasingly important to identify an agenda for future research around human-computer integration.

A common way of identifying "generations" of computing is to consider the ratio of users interacting with computers over time. We began with the one machine/many users paradigm of the Mainframe era, shifting to the one machine/one user paradigm of the PC, and the one user/many machines paradigm of mobiles, to finally the many machines/many users paradigm of today's ubiquitous computing era [1]. However, recent developments suggest a new era where the boundary between machines and users is increasingly blurred as computers are becoming more and more integrated with the users. As such, the trend of HInt emerged, which is a growing interaction paradigm where the computer is closely coupled with the user, physically and conceptually [2]. The emphasis on HInt includes (1) the computers are worn or integrated into the user's body which forms the physical proximity; (2) the computers communicate directly to human senses rather than symbolically which forms the sensory fusion; (3) the computers can influence human task performance which forms a coordinated effort. HInt has an explicit end-goal of merging the human and the computer. Moreover, HInt views ubiquitous computing and good design practices, which suggest more rapid automaticity with and "invisibility" of computers, as waypoints towards a cybernetically integrated future. As such, HInt extends the current paradigm of Human-Computer Interaction (HCI) by introducing physical fusion and partnership with computers. The field of HInt is rapidly expanding, embracing new technologies and incorporating new disciplines. At the same time, the field has matured and is beginning to converge on key questions surrounding technology, self-perception, societal and design implications. We are excited about the potential of the HInt and how it will affect the user experiences with computers.

This seminar began with talks by all the attendees, in which they presented their works in the area, their theoretical perspective that guides their work, a description of their most and least favorite Hint projects, and their expectations for this seminar. After the presentations concluded, the seminar was mainly informed by group discussions, talking about the definition, grand challenges, and the future of human-computer integration. The structure of the seminar was based around theory, design and their intersection. From the start, it was acknowledged that if concerning oneself with technology that is integrated to the human body, not one particular theory will suffice, but rather, that a mix of theories will need to be engaged in, with all their weaknesses and strengths, with the big picture being what we get from studying this.

Figure 1 BioSync.

References

1 Mark Weiser. *The Computer for the 21st Century*. Scientific american 265, 3: 94–105, 1991.
2 Umer Farooq, Jonathan Grudin. *Human-computer integration*. interactions 23, 6: 26–32, 2016.

4 Demo Hour

Joseph La Delfa (RMIT University, Melbourne, joseph@exertiongameslab.org)

On day two of the seminar, participants demonstrated their completed and in-progress works, giving the rest of the researchers a chance to experience some of the human-computer integration concepts and challenges discussed on the opening day.

Jun Nishida from the University of Tsukuba, Japan, presented two experiences which shift the user's perspective such that they are able to embody another person. BioSync [12] is an electronic muscle stimulation (EMS) device that can transform basic movements of the hand from one human to another. "CHILDHOOD" [11] is a visual and haptic perspective-changing experience that consists of two components, an AR experience that shifts the user's vision from eye height to waist height and a mechanical glove that transforms the capabilities of the user's hand into those of a child's hand. These two experiences offer designers a powerful insight into how do design for people with neuromuscular problems and children respectively.

Dag Svanaes from the Norwegian University of Science and Technology demonstrated artificial moveable ears that could be manipulated with a sensor-embedded glove [15]. Through

Figure 2 Artificial ears.

interacting with other people using a pair of artificial ears, one gets an idea of how computers can be used to integrate and augment with body language.

Joseph La Delfa from RMIT University Australia, demonstrated a Tai Chi drone experience [3], where participants were asked to wave their hands from left to right in time with a drone which moved around their body. The experience is designed to question how to design physical objects that can facilitate awareness of body sensations to achieve a relaxed and focussed state.

Joe Marshall from University of Nottingham United Kingdom, presented a perspective-shifting experience which aimed to mimic the kind of disorientation experienced by an infant whilst learning how to coordinate their vision with their movement. The experience highlighted the effort required to seamlessly integrate our natural sensors with our movement and questions what is required to bring an artificial sensor to the same level of integration.

Kai Kunze from Keio University, Japan, demonstrated a pair of glasses [10] capable of tracking facial expression in everyday scenarios. The device raised questions about how this technology could change the relationship between human and computer now that the computer is capable to constantly track and understand facial expressions. Kai also demonstrated a virtual reality juggling game, designed to teach the coordinated movements required to keep multiple objects in the air. The experience was able to "reduce gravity" to give the participant more time to coordinate their movements, raising questions about how computers can dynamically adjust the sensory input to reduce the learning curve of a new skill or augmented ability.

Suranga Nanayakkara from the University of Auckland presented the FingerReader 2.0 [1], a device which embeds a camera into a wearable ring, giving people with a visual disability assistance in identifying objects, text and color by pointing at them. The technology raised questions around the social and phenomenological implications of pointing to see.

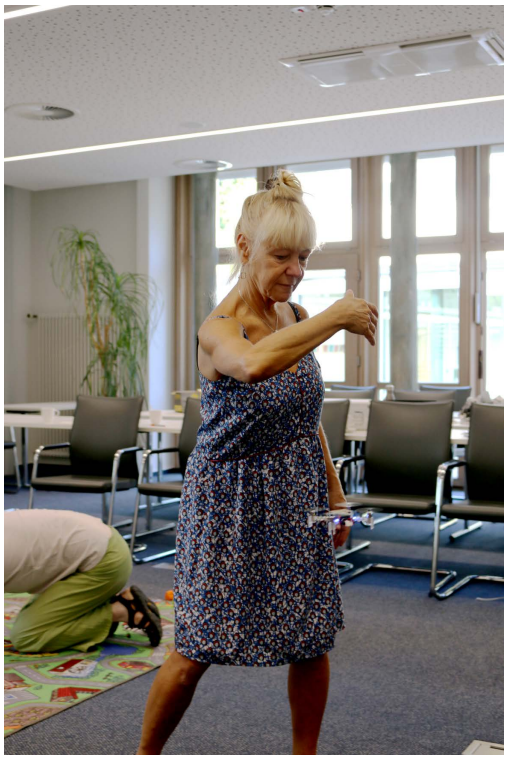

Figure 3 Tai Chi drone.

Caitlyn Seim from Georgia Institute of Technology presented a haptic glove capable of vibrating each finger individually [13]. When worn for long periods of time, the glove can passively teach the user skills such as how to play the piano, read braille, and use a new keyboard. The technology spawned conversations around the possibility of passively learning more complex skills.

The session was important to the seminar as it gave participants a somatic appreciation of the concepts and challenges that were being explored in the other sessions. Furthermore, it was fun and engaging, serving as a welcoming break from the more demanding theoretical discussion.

5 Definition of Human-Computer Integration

Since Human-Computer Integration is an emerging field, a synthesis definition is in demand. During the seminar, all the participants were divided into five groups to discuss the definition of HInt and then shared the results.

Farooq and Grudin's definition of human-computer integration [2] provides a structured description of potential ways in which digital technology and humans could be integrated. This definition suggests that computers and humans can form a partnership to facilitate integration. After the group discussions, we expanded this definition and identified two types of human-computer integration that go beyond the previous definition: (1) the significant reduction in size of sensors and effectors that enables the fusion of technology with the human sensory and motor systems; and (2) massive knowledge bases, networked infrastructure, and intelligent systems that enable a human-computer partnership or symbiosis.

Figure 4 Having the experience of an infant.

We believe that fusion is a type of HInt, which refers to systems that extend the experienced human body. To facilitate the experience of fusion, technologies are merged with the body (e.g., implanted sensors [7], ingested pills [8], and epidermal electronics), extending or manipulating the body [15], or stimulating the senses (as by actuators or augmented reality [9, 13]). Moreover, these technologies might access nearly-limitless information databases and extend our minds. For fusion to occur, interactions with systems and agents must afford direct mediation, perception, and communication. With fusion, we do not command our computers to act, we just act. We do not need to interpret the computer's feedback while we have an embodied understanding of it.

The other type of integration is symbiosis. We believe that as the power of the computer increases, the collaboration between human and computers will occur ultimately and there will be a power balance between what the human can achieve and what the computer can do. We call this type of integration symbiosis. The form of the system may vary, from systems that collaborate in creative tasks [2, 5] to brain implants that selectively trigger memory [4]. The novelty is not that there are software agents, or that the agents are smart. The novelty is that the process is truly shared between system and user as the two act in concert. An agent that has access to and an understanding of the context within which the human operates can adapt its behavior accordingly. In doing so, symbiosis can occur.

6 Grand Challenges of Human-Computer Integration

Shneiderman et al. [14] suggest that HCI as a field needs "grand challenges" to steer the direction of future research, design and commercial development. We believe identifying the grand challenges of HInt could help researchers and practitioners in this filed (1) identify

Figure 5 VR juggling.

current knowledge, capabilities and areas of opportunity where they can contribute; (2) situate their work within the larger HInt research agenda; (3) allow policy makers to better understand the HInt community, state-of-the-art technology and research, and potential applications. After the group discussions, we presented four sets of grand challenges of HInt.

First, the challenge of technology needs to be considered. For example, a close integration between computers and the human body may require the computers to feel and behave like parts of the human body. As such, the materials for the HInt systems need to be considered: it might be beneficial to be biocompatible, miniaturized and deformable. Moreover, the integrated computers need to support the wide range of shapes and sizes of the human body. Also, energy management and harvest is an interesting challenge for HInt devices.

Second, there are some challenges highlighting the identity and behaviors around HInt systems. For example, the HInt systems might change our perceptions of ourselves since such technologies have the potential to change our body schema by enhancing our sensory system and extending our capabilities.

Third, HInt systems may affect people's lives and society. For example, the influence of "digital divide" might be amplified by HInt systems. If areas of public space become designed for people with HInt systems that extend their capabilities, does the sensory-divide created by this design exclude people who cannot afford the augmentation?

Fourth, there are challenges for HInt in the field of interaction design. One of the challenges might be applying novel technologies to develop common understandings and tools for designing, developing, refining, testing and evaluating HInt systems. It is important for research to tackle the issues of robustness and practicality that enable HInt systems to be integrated into our daily lives. In conclusion, we believe the challenges we identified need to be addressed in the future to reap the full benefits of HInt.

Figure 6 FingerReader 2.0.

7 Superheros and Science Fiction

Martin Weigel (Honda Research Europe – Offenbach, DE)
Zhuying Li (RMIT University – Melbourne, AU)
Tom Erickson (Minneapolis, US)

In this session, we spent half an hour exploring the realm of superheros and science fiction to better understand the design space of Human-Computer Integration. The idea was to collect information about the special abilities of these characters as a group exercise and think about how well the abilities are integrated from the perspective of the mind, body and total integration. Groups of four got two characters handed out as a paper card. After discussing and rating these characters, they created a slide comparing them. The slides were presented by the groups to the seminar. The cards were put onto the body-mind integration scale to identify clusters and find empty spaces. As such, this session helped to identify new themes, which were later extracted from the participants' comments in a break-out session.

Based on the results of the discussion, we further created a matrix consisting of values and attributes of HInt. Attributes are aligned to values and can align to more than one value. We went through all the keywords which were used by the participants to describe the integration of the superheros. Then we classified these keywords as value or attribute. If it is an attribute, we identified what value is it associated with. We concluded values including safety, agency, extensions (something that extends the capability of an ecosystem), ethics/moral, aesthetics, and so on. Values determine how the attributes operate. We believe the values and attributes we concluded could also help make up the design space of HInt systems.

Figure 7 Haptic glove.

8 Methods of Tools of HInt

Wendy Ju (Cornell Tech, New York)
Zhuying Li (RMIT University, Melbourne)
Tom Erickson (Minneapolis, US)

We realize that the emerging field of HInt needs methods and tools. The seminar participants discussed the potential methods and tools of HInt. The assumption here assumes that the Human-Computer Integration curriculum follows an HCI 101 course. The ways that the HInt curriculum goes beyond HCI fundamentals is that: (1) HInt has a greater degree of integration with the user's body; (2) HInt systems are more "analogy" while traditional HCI projects are mainly on/off systems; (3) HInt systems are smaller and have more complex signals; (4) HInt systems has a different time scale compared to HCI systems; and (5) HInt brings about simulation and reality alignment issues.

We realized that interaction quality is important in HInt systems. To address this, In-situ/ simulation design and development including participatory design might be needed. After the prototype has been developed, an always-on test that continuous for 1-2 days is also necessary. Methods of instrumenting environments might also be applied to testing HInt systems.

When it comes to tools for HInt, hardware like Arduino and coding language such as Python might be utilized. On the mechanical side, it would be good to have stuff on adhesives, connectors, straps/sleeves/harnesses, etc. Other knowledge about signal processing, machine learning, control systems, eye tracking and VR/AR is also beneficial for HInt. Moreover, there might be an interest in getting into biochemical or pharmacological manipulations.

Figure 8 Discussing the definition of HInt in groups.

Certain design curricula is also needed for developing HInt systems. For example, designers should consider the perspective (1st, 2nd, and 3rd person perspectives) they take. Infrastrucuring might help designers take a socio-technical perspective on things. Other design knowledge from ethnography, critical design, empathic/experiential design could help the HInt system design. Designers of HInt might also need to think about the 2nd-order and 3rd-order effects and the longer term effects of design on life, people and societies. Simulation, gaming and forecasting might be used to design to anticipate these later scale effects.

9 Speculative Futures of HInt: Investigating User Scenarios

Ti Hoang (RMIT University, Melbourne, ti@exertiongameslab.org)
Pattie Maes (MIT, Cambridge, pattie@media.mit.edu)

To frame the emerging field of human-computer integration (HInt) we workshopped the concept of speculative futures for HInt to identify user cases and issues for designers of future HInt scenarios. Health and education were two specific applications explored in-depth. In each scenario, we identified the potential end-users of the Hint technology, the possible scenarios where HInt could occur, and finally developed thumbnail sketches of possible scenarios. Through this, we were able to identify fifteen key themes that represented either barriers to entry or under-explored directions for designers of HInt. Those key themes included: privacy, loss of control, security and safety, technical malfunction, graceful degradation, dependency, unequal access, disparity, techno tribalism, simulations to increase empathy, discontinuation of services, expectations of the technology, addiction to technology, psychological and social impact, and environmental impact and sustainability.

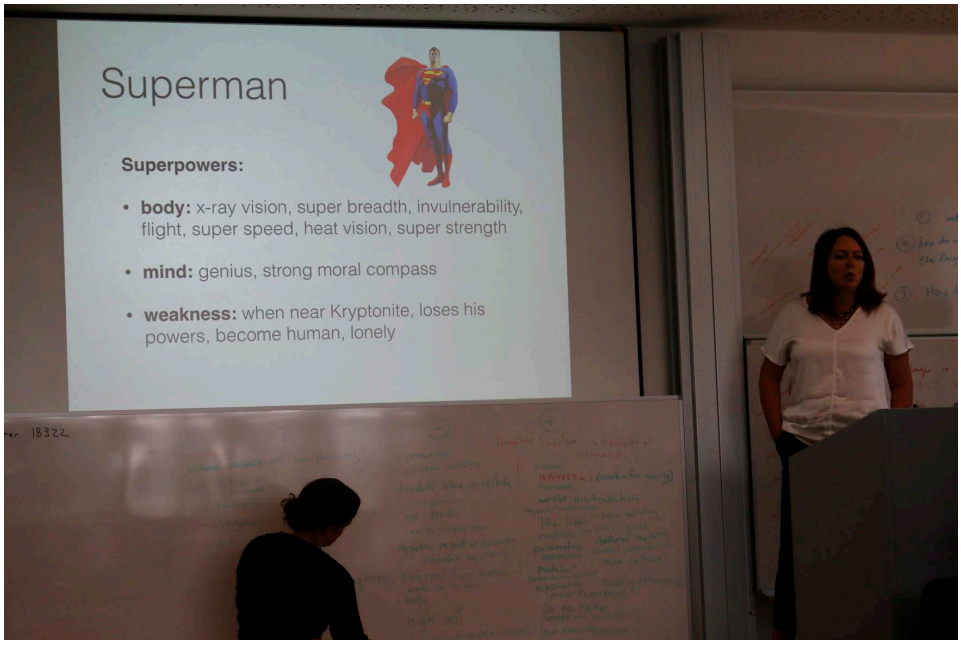

Figure 9 Discussing the level of human-computer integration of superheroes.

In attempting to frame the future of human-computer integration we offer the name HInt as a way to describe and differentiate human-computer integration from the previous focus of human-computer interaction. We classified HInt technology as either off-body, on-body, or in-body. We defined "off-body" as technology which is situated in the environment around the body and does not physically attached to the body, "on-body" was defined as technology that exists on the surface of the body such as wearables or hand-held devices and which can be separated from the body, and "in-body" as technology which exists internally within the body such as ingestible devices. We also draw inspiration from comic book characters with superpowers as a fictional representation of HInt for the purpose of suspending disbelief and encouraging creative thinking to consider alternative futures of HInt.

To identify future possible scenarios for HInt we undertook an informal brainstorming and discussion session which focused on what we believed would be the future of HInt. The discussion involved four phases not too dissimilar to the ideation and concept generation process used within design disciplines: brainstorming ideas to generate a set of possible directions for Hint, identifying and selecting specific applications from the set of possibilities for further investigation, and generating user scenarios for those applications. Designing and developing mock-ups of user scenarios was outside of the scope of the activity, and instead, the intent was to identify themes to assist designers in developing future HInt systems.

The initial brainstorming phase resulted in the mapping out of the broad philosophical HInt questions: What is the HCI discipline in a decade? What will be new approaches to designing HInt? Will these systems be distributed? What are the software agents? How can HInt facilitate social interactions or assist in unifying the self? What are exemplars of HInt? And is the integration of technology with the body even a path we should be undertaking? To explore possible HInt health scenarios we first identified the users of the system, and they included patients, extended family, or health practitioners. One possible future HInt scenario re-imagined how practitioners might practice healthcare. It

Figure 10 Lay on the floor and feel the 1st person perspective of live experience.

was suggested that practitioners could use technology such as virtual reality to walk through the virtual body of a patient. Another suggestion proposed assistive technology could be developed to guide users in learning to eat well, prepare for surgery, or experience what others experience. The HInt user scenarios were classified into three areas: transformative powers, restorative powers, and superpowers. The use of the term "powers" was a playful prompt for re-imagining the outcomes of HInt and draws from the notion that comic book characters with superpowers are fictional representations of HInt. Our second scenario investigated educational Hints. The users of HInt learning were students and human tutors, and future HInt scenarios might include downloadable skills, technology to enhance the human senses, and deep brain stimulation. For on-body HInt scenarios, the group generated three possible scenarios: deep brain stimulation, sensory enhancement, and mobile learning using AR. This scenario generated concepts such as wearable clothing, or a hand-held stick, that can assist in translating languages.

While exploring user scenarios for each application, we also identified how this technology might break down. Concerns included: privacy, loss of control, security and safety, technical malfunctions, discontinued software, dependency and addiction, environmental sustainability, psychological and social impact, disparity brought about from the cost of technology and policy making. More intriguing breakdowns and perhaps specific to HInt technology came after the initial list of issues was identified. Concerns of techno tribalism where users could potentially become technologically isolated due to the brand or manufacturer of HInt technology that they choose; body-mining of technology where the value of HInt technology could result in crime; and heirloom HInt technology where people might associate the sentimental value to technology in unexpected ways.

In summary, the discussion highlighted the need to investigate where future HInt scenarios might be needed, and that designers of HInt should also be aware of the repercussions of this technology. Design approaches which allow people to re-imagine future HInt scenarios could help with developing more detailed narratives which assist to move the discussion towards richer visual and physical representations of speculative future HInt scenarios.

Acknowledgement. We thank Dagstuhl for their extensive support, and all the participants who contributed.

References

1. Roger Boldu; Alexandru Dancu; Denys J. C. Matthies; Thisum Buddhika; Shamane Siriwardhana; Suranga Nanayakkara. *FingerReader2.0: Designing and Evaluating a Wearable Finger-Worn Camera to Assist People with Visual Impairments while Shopping*. In Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 2, 3, 2018.
2. Mason Bretan; Gil Weinberg. *Integrating the Cognitive with the Physical: Musical Path Planning for an Improvising Robot*. In AAAI: 4371–4377, 2017.
3. Joseph La Delfa; Robert Jarvis; Rohit Ashok Khot; Florian Mueller. *Tai Chi In The Clouds: Using Micro UAV's To Support Tai Chi Practice*. In Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts: 513–519, 2018.
4. Youssef Ezzyat; Paul A Wanda; Deborah F. Levy; Allison Kadel; Ada Aka, Isaac Pedisich; Michael R. Sperling; Ashwini D. Sharan; Bradley C. Lega; Alexis Burks. *Closed-loop stimulation of temporal cortex rescues functional networks and improves memory*. Nature communications 9, 1: 365, 2018.
5. Ashok K. Goel; Spencer Rugaber. *Interactive meta-reasoning: Towards a CAD-like environment for designing game-playing agents*. In Computational creativity research: Towards creative machines: 347–370, 2015.
6. Jonathan Grudin. *Computer-supported cooperative work: History and focus*. Computer 27, 5: 19– 26, 1994.
7. Christian Holz; Tovi Grossman; George Fitzmaurice; Anne Agur. *Implanted user interfaces*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: 503–512, 2012.
8. Zhuying Li; Rakesh Patibanda; Felix Brandmueller; Wei Wang; Kyle Berean; Stefan Greuter; Florian Mueller. *The Guts Game: Towards Designing Ingestible Games*. In The Annual Symposium on Computer-Human Interaction in Play Extended Abstracts (CHI PLAY '18): 271–283, 2018.
9. Pedro Lopes; Ryan Lavoie; Rishi Faldu; Nick Aquino; Jason Barron; Mohamed Kante; Basel Magfory; Waleed Meleis. *Eye-Controed Robotic Feeding Arm Technology*. 2012.
10. Katsutoshi Masai; Kai Kunze; Yuta Sugiura; Masa Ogata; Masahiko Inami; Maki Sugimoto. *Evaluation of Facial Expression Recognition by a Smart Eyewear for Facial Direction Changes, Repeatability, and Positional Drift*. ACM Transactions on Interactive Intelligent Systems 7, 4, 2017.
11. Jun Nishida; Hikaru Takatori; Kosuke Sato; Kenji Suzuki. *CHILDHOOD: Wearable Suit for Augmented Child Experience*. In Proceedings of the 2015 Virtual Reality International Conference, 2015.
12. Jun Nishida; Kenji Suzuki. *bioSync: A Paired Wearable Device for Blending Kinesthetic Experience*. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems: 3316–3327, 2017.
13. Caitlyn Seim; John Chandler; Kayla DesPortes; Siddharth Dhingra; Miru Park; Thad Starner. *Passive haptic learning of Braille typing*. In Proceedings of the 2014 ACM International Symposium on Wearable Computers: 111–118, 2014.
14. Ben Shneiderman; Catherine Plaisant; Maxine Cohen; Steven Jacobs; Niklas Elmqvist; Nicholoas Diakopoulos. *Grand challenges for HCI researchers*. interactions 23, 5: 24–25, 2016.

15 Dag Svanaes; Martin Solheim. *Wag Your Tail and Flap Your Ears: The Kinesthetic User Experience of Extending Your Body*. In Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems: 3778–3779, 2016.

10 Overview of Talks

10.1 Tom Erickson

Tom Erickson (Minneapolis, US)

After a short introduction characterizing myself and my approach, I set out three prejudices I bring to the workshop:

1. AI is stupid and brittle and is not likely to get much better.
2. Affect recognition is crude, and is not likely to get much better.
3. and therefore "human computer partnership" – a phrase often used to characterize the future of human computer interaction – is not likely to get much better, because partnership requires high cognitive and emotional intelligence.

That said, the concept of partnership need not be restricted to reciprocal relationships among co-equal partners. To enquire more deeply into partnership, I discuss a paper on "Human-Sheepdog Distributed Cognitive Systems," which is interesting because it illustrates partnership between three entities that are deeply unequal. The paper analyzes how the 'partnership' works

- the shepherd providing the plan and large-scale sensing of where the sheep are with respect to the goal, and signaling the dog on where to move the sheep
- the dog providing more acute but more local sensing of the sheep (with no awareness of the goal), and acting to move and control the sheep
- and the sheep sensing the dog, who is a predator that they fear, and trying to keep a comfortable distance away (thus, being herded)

It is interesting to note that the three entities have entirely different 'views' of what is going on:

- the shepherd is participating in a sheep-herd trial
- the dog is enacting predator routines[1], and responding to signals from its trainer
- and the sheep are responding to a potential predator.

This is an interesting example of an ecosystem of coordinated but not really cooperating entities acting in a coherent fashion. Perhaps applying this sort of distributed cognition analysis to human-computer partnership might be a way to grapple with the deeply unequal capabilities of humans and computers.

[1] The dog's predator routines have been subverted through nurture – sheepdog puppies are raised with sheep, so that while the sheep trigger their predator routines they also identify them as their 'pack' and do not attack them – and training.

10.2 Elizabeth Gerber

Elizabeth Gerber (Northwestern University – Evanston, US)

It takes a village to raise a child. Parents provide housing and moral guidance. Teachers provide instruction, etc. I presented the idea of understanding how technology raises ideas. Technologies may support social exchange – social media, online communities, live video chats, instructional videos, information exchange – crowdsourcing, or financial exchange – crowdfunding. My research mission is to use computation to advance social prosperity. My work is driven by Hutchins and LaTour's perspective of distributed cognition where cognitive processes are distributed across the members of a social group and technology facilitates exchange. I conduct this research in the Delta Lab at Northwestern University designed to bring together computer scientists, learning scientists, and organizational behavioralists to study this topic. I was formally trained as a designer and organizational behavior scholar at Stanford University.

10.3 Steve Greenspan

Steven Greenspan (CA Labs, US)

Thank you for inviting me to the Human-Computer Integration workshop (Dagstuhl Aug 6-10, 2018). My name is Steve Greenspan – a quick background: My PhD is in Cognitive Psychology and this was followed by postdocs in user experience under Don Norman and speech perception under David Pisoni. As a researcher at AT&T Labs, I led a research program in 2002 called Air Graffiti, in which a head-ups display presented audiovisual information based on the location of the user – when users entered an office with an occupant they might receive a Zork message "you've entered a room and there is a wizard at the desk", or when they passed an office co-worker's door they might hear something about the worker's latest activity – whatever the worker wished to post to passers-by.

More recently at CA Technologies, I have been involved in research projects on differential privacy, ML-facilitated visual analytics, and AI ethics. Of relevance to this workshop, I am currently working on Cobotics and IoT data trustworthiness. In the Cobotics work we are researching how teams of humans and teams of robots cooperate and coordinate work. The work on data trustworthiness is focused on how to make good decisions when the data from IoT sensors is not completely trustworthy.

There are many challenges to human-computer integration.

1. What stimuli can be mapped to what senses? Can we use synthesia to map new senses (e.g., mapping magnetic field strength to auras around objects in a person's visual field)?
2. What are the limits on the number of inputs that can be cognitively integrated?
3. Preventing Attack Vectors (security by design)
   - What are the new attack vectors? Perceptual sensitivities, illusions, maladaptions, fake affordances.
4. Ethical & Legal implications: Privacy & Responsibility

- Who is responsible for unintended and intended harm: The manufacturers? The service providers? Or the user? Under what scenarios?

We must be concerned not only with how the user integrates with new sensor and motor devices, but also with how these human-machine systems coordinate with other human-machine systems and with society.

10.4 Playful Interaction

Stefan Greuter (Deakin University – Melbourne, AU)

My presentation on "Playful Interaction" introduced previous work including the design of an interactive tablet based learning game to teach Occupational Health and Safety on Australian construction sites; Virtual Reality projects for architectural and cultural heritage visualisation, games for a full-dome system and playful interactions or industrial machines and laboratory environments. Play is important as it allows us to learn, use our creativity and imagination, improve our dexterity and physical, cognitive and emotional strength. Therefore, as our interaction with intelligent and highly integrated and automated systems progresses we need to focus on human factors and values such as our playful human nature.

10.5 Human-Computer Integration

Jonathan Grudin (Microsoft Research – Redmond, US)

In the half a century since the first graphical user interface supplemented program statements and command lines, we have shifted from software that waits for the next command to software that is continually acting on behalf of people, with or without their awareness. These human-computer partnerships, or symbioses, or integration, alongside new hardware and pervasive networking, have opened a range of development possibilities and research opportunities or even imperatives. My recent work on interactive conversational agents has revealed to me how difficult this will be. Yet it is also a great opportunity—in order to build software that simulates aspects of being human, we must understand better how human beings think and act, and when we understand that, we can support it better whether we are building conversational agents or other software and hardware.

10.6 Re-imagining the digital climbing experience

Ti Hoang (RMIT University – Melbourne, AU)

To design novel exertion games there needs to be a shift in our understanding of the ways in which interactions can occur between humans and technology. At the centre of this paradigm

shift is the human body, specifically how we view the human body and the role it plays within Human-Computer Interactions (HCI).

Advances in interactive technology and their capacity to be tightly coupled with the body are opening new opportunities to explore how designers within HCI might design future exertion games that can strengthen the relationship between the body, the mind, and the environment. Rock climbing is an ideal sport for converting into a computationally-augmented exertion game because the whole body and the mind are intrinsically engaged when performing in this sport. The development of climbing gyms has also assisted in making the sport more accessible to the general public by providing a relatively safe environment for people to experience rock climbing. This increase in popularity has led to research into augmented climbing experiences. Past augmented climbing experiences have focused on: gamifying the sport by projecting digital games onto climbing walls; simulating real-world climbing experiences by creating virtual reality environments; or providing training assistance. They are also predominantly technology-driven approaches to developing digital climbing experiences, and have focused on engagement with the climbing wall. This suggests there is potential to explore the other components, such as the climber, in this experience.

This research begins by exploring rock climbing but is not limited to this sport, and has further applications for understanding how we design for exertion games. There are two emerging and distinct ways in which exertion games are being augmented, and that is either: augmenting the environment, or augmenting the body. My own research will explore how our bodies can be both physically and mentally engaged within exertion games. It will investigate ways of coupling emerging technology with the human body to extend the body's capacity to perform and experience play.

10.7 Misahiko Inami

Masahiko Inami (University of Tokyo, JP)

Masahiko Inami is a professor from University of Tokyo, Japan. His research interest is in human I/O enhancement technologies including bioengineering, HCI and robotics. He received BE and MS degrees in bioengineering from the Tokyo Institute of Technology and PhD in 1999 from the University of Tokyo. His research exploits all five senses for interaction. He proposed the concept of haptic augmented reality via projects such as SmartTools and SmartFinger. His team has archived several improvements that use multi/cross modal interfaces for enhancing human I/O. They include Transparent Cockpit, Stop-Motion Goggle, Galvanic Vestibular Stimulation, JINS MEME (electrooculography (EOG)-based smart glasses) and Superhuman Sports.

Professor Inami believes that today's Human-Computer Interaction (HCI) systems include virtual/augmented reality are limited, and exploit only visual and auditory sensations. However, in daily life, we exploit a variety of input and output modalities, and modalities that involve contact with our bodies can dramatically affect our ability to experience and express ourselves in physical and virtual worlds. Using modern physiological understandings of sensation and perception, emerging electronic devices, and agile computational methods, we now have an opportunity to design a new generation of "Human-Computer Integrated" systems.

10.8 Integration through Interaction

Wendy Ju (Cornell Tech – New York, US)

I make the case that Integration occurs through Interaction. In my research looking at how to design interactions with automation, I have discovered that often what feels natural and obvious as an interaction is different from what makes sense, cognitively or logically, as a model. This is because people engage with the world phenomenologically in the moment of interaction. Hence, in my experiment, we often perform design intervention experiments in which we elicit interaction patterns from people in-situ; often we use wizard of oz techniques wherein the wizard's instinct for interaction is as much a part of the experiment as the users' reactions. I argue that designing an interaction is like designing a conversation-you can't do it by yourself. It goes in a lot of directions; it depends a lot on context. Though I design interactions with robots and autonomous cars, I most frequently look to the writings of linguists such as Susan Brennan and Herb park for inspiration and explanation of how people communicate non-symbolically. My early work on The Design of Implicit Interaction, which looks at interactions where people are not consciously inputting information or seeking information out, illustrated how implicit commands and implicit displays change interaction patterns so that they are fundamentally different than traditional HCI interactions wherein explicit commands and displays are the norm. The advent of automation and integration make the question of how to design such interactions crucial to moving HCIntegration from technological possibility to reality.

10.9 Kai Kunze

Kai Kunze (Keio University – Yokohama, JP)

I work on technology to understand ourselves better. I'm an Associate Professor at the Keio Graduate School of Media Design, Keio University, Japan. I love science, hacking, making and playing with tech. In my research, I combine design and technology to make human experiences sharable to capture and exchange abilities, ultimately to amplify human senses. My overall goal: I want to give people a toolset to improve their physical and cognitive skills applying technology and design to enhance attention, comprehension, memory and ultimately decision making.

10.10 Joseph La Delfa

Joseph La Delfa (RMIT University – Melbourne, AU)

I introduced the octopus as a potential model of human computer integration, coving some basic concepts around cognition, proprioception and 'dependency'. Specifically the idea that the successful integration of the octopus's smaller ganglia and its larger brian are partially

due to the co-dependency. I then followed with an example of the implications of wearing a garment that makes social decision on your behalf.

10.11 Play with Human-computer Integration

Zhuying Li (RMIT University – Melbourne, AU)

I am a PhD student with a background of Engineering and Game Design. I am interested in the future of play in the era of human-computer integration. I believe with the emergence of augmented-human technologies, more symbiotic human-computer interfaces will be developed to help create novel play experiences. I further explained this in my presentation by introducing two of my projects Guts Game and HeatCraft, which use ingestible sensors to facilitate playful experiences.

10.12 Devices that Overlap with the User's Body

Pedro Lopes (Hasso-Plattner-Institut – Potsdam, DE)

How can interactive devices connect with users in the most immediate and intimate way? This question has driven interactive computing for decades. If we think back to the early days of computing, user and device were quite distant, often located in separate rooms. Then, in the '70s, personal computers "moved in" with users. In the '90s, mobile devices moved computing into users' pockets. More recently, wearables brought computing into constant physical contact with the user's skin. These transitions proved to be useful: moving closer to users and spending more time with them allowed devices to perceive more of the user, allowing devices to act more personal. The main question that drives my research is: what is the next logical step? How can computing devices become even more personal?

Some researchers argue that the next generation of interactive devices will move past the user's skin, and be directly implanted inside the user's body. This has already happened in that we have pacemakers, insulin pumps, etc. However, I argue that what we see is not devices moving towards the inside of the user's body but towards the "interface" of the user's body they need to address in order to perform their function.

This idea holds the key to more immediate and personal communication between device and user. The question is how to increase this immediacy? My approach is to create devices that intentionally borrow parts of the user's body for input and output, rather than adding more technology to the body. I call this concept "devices that overlap with the user's body". I'll demonstrate my work in which I explored one specific flavor of such devices, i.e., devices that borrow the user's muscles.

In my research I create computing devices that interact with the user by reading and controlling muscle activity. My devices are based on medical-grade signal generators and electrodes attached to the user's skin that send electrical impulses to the user's muscles; these impulses then cause the user's muscles to contract. While electrical muscle stimulation (EMS) devices have been used to regenerate lost motor functions in rehabilitation medicine since

the ’60s, during my PhD I explored EMS as a means for creating interactive systems. My devices form two main categories: (1) Devices that allow users eyes-free access to information by means of their proprioceptive sense, such as a variable, a tool, or a plot. (2) Devices that increase immersion in virtual reality by simulating large forces, such as wind, physical impact, or walls and heavy objects.

10.13 Cognitive Enhancement

Pattie Maes (MIT – Cambridge, US)

Most digital tools and applications today are designed for interaction with the conscious, logical, rational-thinking part of a user, with communication between the person and the technology being of a purely symbolic nature, i.e. words and pictures, and requiring the user’s complete attention. But people are more than rational, slow thinkers. A lot of our thinking and behavior is automatic, instinctive and emotional and is heavily influenced by the senses. I advocate for designing technology for the “whole” person, i.e. for both the rational and instinctive parts. I would like to encourage the HCI community to read more psychology and neuroscience, learn about the intricacies and oddities of the human brain and behavior, exploit those in new types of integrated-in-the-body interfaces that make use of non-symbolic means of interaction with the user. I believe such interfaces will be more effective in influencing and impacting users at a deep level so as to support memory, attention, learning and behavior change.

10.14 Joe Marshall

Joseph Marshall (University of Nottingham, GB)

I presented a model of embodied systems in which users experience both digital and real-world physical sensory stimulation. Taking a holistic multi-sensory perspective all current systems expose users to a mixture of real and digital sensory experiences. Our model of these digital and physical sensory stimulation mixtures considers 2 things – firstly how each sense is digitally or physically stimulated, and the congruence of digital stimulation, which relates to whether the digital stimulation is consistent with physical stimulation or whether it is imperceptibly or perceptibly inconsistent. I showed 3 examples of deliberate and perceptible incongruence – VR Playground, a playground swing with VR, where the user’s virtual motion is driven by the swing, but mapped to be very different, and Balance Ninja and AR fighter, two games which present inconsistent balance cues by affecting the user’s vestibular system (Balance Ninja), and rotating their vision (AR Fighter).

10.15 Longterm Self Tracking

Jochen Meyer (OFFIS – Oldenburg, DE)

Technology enables us to monitor our behavior and vital parameters. Since some years consumer products are available that can be used in daily life, by laypersons, and over years, decades, and ultimately lifelong.

This opens tremendous opportunities for supporting healthy living such as identification of slow changes and making decisions about future health behavior.

In such scenarios the user has a central role, acting not just as a consumer of services that are provided by a technical system, but also as a producer of data that is input to the system. These two roles may have competing requirements: As a consumer the user usually wants highest-quality services, which usually require high-quality data. As a producer, however, the user's interest more often is to reduce effort of data collection as much as possible, resulting in low-quality data with gaps and holes, which in turn limits the possibilities of using the data in applications.

This results in a mutual dependency between the application and the data: the applications defines requirements to the quality of data to fulfill the requested services as given by the user as a consumer; on the other hand the available data, as given by the user's role as a producer, defines the possibilities of the application for providing services. Balancing these two views is a key challenge for long-term applications.

10.16 Body-Computer Integration

Florian Mueller (RMIT University – Melbourne, AU)

URL http://exertiongameslab.org

I propose that one part of a human-computer integration future is body-computer integration. This is grounded in philosophy that argues for the importance of human values that ultimately ends in the good life. I illustrate this thinking by presenting recent work from the Exertion Games Lab, including a traffic-light aware eBike, a singing ice-cream and on-body robotic arms for social dining experiences.

10.17 Suranga Nanayakkara

Suranga Nanayakkara (University of Auckland, NZ)

URL www.ahlab.org

My 7-8 year experience working with deaf children made me realize that sensory impairment has nothing to do with intellectual ability. For instance, these deaf children were able to communicate over much longer distances with sign language and make beautiful computer graphics. In fact, they have such a developed special skill that I felt like the odd man.

Therefore, I believe in Human-Human Integration where technology act as an enabler to connect different communities with different abilities, helping people do things that they think they could not. At Augmented Human Lab (www.ahlab.org), we are taking some initial steps towards this.

10.18 Jun Nishida

Jun Nishida (University of Tsukuba, JP)

My presentation was about "HYPERSPECTIVE: Egocentric Perspectives by Wearable I/O Devices to Understand, Communicate and Cooperate with People." I introduced wearable devices to change the sense of body into an another person including 1) a wearable AR device that transforms the height of perspective into that of a child, 2) hand exoskeletons to change hand dimensions into that of a child, and 3) a pair of wearable muscle I/O device to share kinesthetic experience among people. These devices allow people to experience one's perspective in an embodied and egocentric manner. I believe that these hyperspective experiences would not only enhance our embodied knowledge but also change people's behaviour as well. When integrating humans and computers, it would be very important to preserve user's egocentricity and agency, because HInt devices are capable of deeply interfering user's actions, and having a full control of a user. To empower users with wearable systems, new technologies to understand a user's intentions, and to blend a device intervention and a user's voluntary are required. I hope my work encourages studies in this topic, and provide design implications to future researches in HInt.

10.19 Multisensory Experiences

Marianna Obrist (University of Sussex – Brighton, GB)

We experience the world around us through all our senses, but the way technology is designed is often limited to the stimulation of our sense of vision and hearing. In my presentation I argue for the integration of touch, taste, and smell into interactive technology. Especially the chemical senses, taste and smell, are under-exploited and yet today our understanding on those senses is more advanced than ever due to breakthroughs in psychology, neuroscience, and sensory science. Therefore we can now more actively design multisensory experiences in the future integration of human and technology. The benefit of doing that integration lies in the use of smell and taste for various application scenarios including training, education, therapy, entertainment, etc. – exploiting the strong link of those sensory stimuli to emotions and memory. My personal vision is to establish a systematic understanding of what and how to design for tactile, gustatory, and olfactory experiences in human-computer interaction (HCI), for life and experiences on Earth and beyond!

10.20 Harald Reiterer

Harald Reiterer (Universität Konstanz, DE)

Prof. Dr. Mag. Harald Reiterer holds a Magister (Mag.) degree (M.Sc. equivalent) from the University of Vienna in Computer Science and Economics. He defended his Ph.D. thesis in Computer Science at the University Vienna, Austria in 1991. In 1995 the University of Vienna conferred him the venia legendi (Habilitation) in Human-Computer Interaction. Prior to his appointment as full professor at the Computer and Information Science Department of the University of Konstanz in 2009, he was associate professor at the Department of Computer and Information Science of the University of Konstanz (1997-2009), assistant professor at the Department of Computer Science at the University of Vienna (1995-1997), and senior researcher at the Fraunhofer Institute for Applied Information Technology (1990-1995) in Bonn, Germany. His main research interests include different fields of Human-Computer Interaction, like Interaction Design, Usability Engineering, and Information Visualization.

10.21 Thecla Schiphorst

Thecla Schiphorst (Simon Fraser University – Surrey, CA)

Thecla Schiphorst is Associate Director and Associate Professor in the School of Interactive Arts and Technology at Simon Fraser University in Vancouver, Canada. Her background in dance and computing form the basis for her research in embodied interaction, focusing on movement knowledge representation, tangible and wearable technologies, media and digital art, and the aesthetics of interaction. Her research goal is to expand the practical application of embodied theory within Human Computer Interaction. She is a member of the original design team that developed Life Forms, the computer compositional tool for choreography, and collaborated with Merce Cunningham from 1990 to 2005 supporting his creation of new dance with the computer. Thecla has an Interdisciplinary MA under special arrangements in Computing Science and Dance from Simon Fraser University (1993), and a Ph.D. (2008) from the School of Computing at the University of Plymouth.

10.22 Caitlyn Seim

Caitlyn E. Seim (Georgia Institute of Technology – Atlanta, US)

My background in engineering and doctoral education in human-centered computing at the Georgia Institute of Technology inform my theories of the upcoming socio-technical shift to more integrated technologies.

Examples of this budding integration include HUDs that inform rapid decision making in order picking, sonic and tactile signals that help regulate respiration and stress, and algorithms that co-create art (images and songs) with human users.

My work on haptics and passive motor learning/rehab is one supporting example of how biosignals will help devices influence users as we become more integrated with technology.

10.23 Jürgen Steimle

Jürgen Steimle (Universität des Saarlandes, DE)

Jürgen Steimle is a full professor at the Department of Computer Science and at the DFG-Cluster of Excellence "Multimodal Computing and Interaction" at Saarland University. He is the head of the Human-Computer Interaction and Interactive Technologies Lab. He is a Senior Researcher at the Max Planck Institute for Informatics. Previously, he was an independent research group leader (W2) at Saarland University and Max Planck Institute for Informatics. From 2012 to 2014, he was a Visiting Assistant Professor and Research Affiliate at Massachusetts Institute of Technology (MIT) Media Lab. He holds a PhD in Computer Science from Darmstadt University of Technology. His research investigates user interfaces which seamlessly integrate digital media with the physical world, in order to enable more effective, expressive, and engaging interactions with computers. His current focus areas include flexible displays and sensor surfaces, on-body interaction, embedded user interfaces, and personal fabrication.

10.24 Designing for and leveraging Active Perception

Paul Strohmeier (University of Copenhagen, DK)

Material properties of the world around us are revealed to us by our interactions with them. It is tempting to think of our environment as having fixed properties which we perceive through passive sensors, but various studies suggest that the pre-conscious mircointeractions between our body and its environment create our subjective experience of the world. This becomes particularly apparent when studying haptic perception. Let us analyze lifting up an object. When holding the object, the fingertips are distorted due to shear stress. This distortion of the fingertips while the object is being lifted leads to a perception of weight [1]. When holding it, there is an interaction between the compression of the fingertip and the corresponding displacement of the fingers through the object. This interaction leads to a perception of compliance [2]. When moving our fingertip over the texture of the object, the interaction between our fingerprints and the materials surface structure causes vibrations. These vibrations are perceived as texture [3]. We experience the world around us through our interactions with the world. This is relevant for HCI as it allows us to provide users with material impressions without recreating the entire material. Rather through studying the sensory modality one wishes to target, one can create the target material by creating tightly coupled feedback loops, simulating the interaction rather than the material properties [4, 5]. This allows us to create perceptions of virtual worlds without needing to recreate the entire world, it also provides us with guidance regarding the design completely new senses and experiences.

References

1 Roland S. Johansson; J. Randall Flanagan. *Coding and use of tactile signals from the fingertips in object manipulation tasks*. Nature reviews. Neuroscience 10, 5: 345–59, 2009.
2 Wouter M. Bergmann Tiest; Astrid M. L. Kappers. *Cues for haptic perception of compliance*. IEEE Transactions on Haptics 2, 4: 189–199, 2009.
3 Sliman Bensmaïa; Mark Hollins. *Pacinian representations of fine surface texture*. Perception & psychophysics 67, 5: 842–854, 2005.
4 Paul Strohmeier; Sebastian Boring; Kasper Hornbæk. *From Pulse Trains to "Coloring with Vibrations": Motion Mappings for Mid-Air Haptic Textures*. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems – CHI'18: 1–12, 2018.
5 Paul Strohmeier; Kasper Hornbæk. *Generating Haptic Textures with a Vibrotactile Actuator*. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems – CHI'17, 4994–5005, 2017.

10.25 Dag Svanæs

Dag Svanaes (NTNU – Trondheim, NO)

Svanæs received his Ph.D. in Human-Computer Interaction (HCI) from NTNU. His research over the last 15 years has been in the fields of HCI and Interaction Design. His main focus has been on user-centered design methods and basic theory of interaction. A common theme is the importance of non-cognitive aspects of human-computer interaction – often called embodied interaction. At a practical level this involves a focus on the physical, bodily and social aspects of interaction. In his research he makes use of role play and low-fidelity prototyping in realistic settings to involve end-users in the design process.

10.26 A New Paradigm Of Human-Computer Interaction: Human-AI Collaboration

Dakuo Wang (IBM T.J. Watson Research Center – Yorktown Heights, US)

We are witnessing an emerging new paradigm of how people interact with computer systems. A few examples include computers we ate into our stomachs, artificial limbs controlled by electrical muscle stimulation (EMS), e-tattoos on our skins, various Augmented Reality/Virtual Reality (AR/VR) systems, auto-piloted cars, and various assistants enabled by Artificial Intelligent (AI) in our phones, at work, and at home. Contrary to the known paradigms, where we know how people use Personal Computers (PCs) with Graphical User Interfaces (GUI), web-based systems, mobile applications (Ubiquitous Computing), social network systems (Social Computing), we have little knowledge about how people interact with these new technologies. We know so little that we do not even have a consensus of the name for this new paradigm, thus the design of such systems is quite opportunistic. In this talk, I will firstly provide an account and a definition for this new paradigm. Then, I propose

the "Human-AI Collaboration" model, which can be used as a guideline for understanding people's interactions with AI systems, and as design principles for designing such systems.

10.27 Cooperative Intelligence

Martin Weigel (Honda Research Europe – Offenbach, DE)

In my short presentation, I briefly introduced myself and my recent Ph.D. work. I also introduced the Honda Research Institute Europe, where I currently work, and our idea of Cooperative Intelligence. This idea is closely related to Human-Computer Integration. Afterwards I presented my personal thoughts on intelligent systems, technical challenges and Human-Computer Integration in general.

10.28 Katrin Wolf

Katrin Wolf (HAW – Hamburg, DE)

Katrin Wolf is a professor for Media Informatics at the Hamburg University of Applied Science in the faculty of Design, Media & Information and before she was a professor for Media Informatics at the BTK, the University of Art and Design in Berlin. She had worked as a postdoctoral researcher in the Human Computer Interaction Group at the University of Stuttgart, where she worked in the meSch project on projected guidance systems and in the RECALL project on lifelogging video navigation.

Her research focus is on human-computer interaction. Particularly, she is interested in touch and gesture-based interaction with augmented environments to foster integrated human-computer interaction. Katrin worked with physiotherapists focusing on ergonomics in gestural interaction design. She is known in the embodied interaction community for her work on microgestures using wearable sensors for gesture detection and exploring implicit interactions as well as explicit gesture interaction for wearable and ubiquitous computing.

PickRing is a wearable sensor that allows seamless interaction with devices through predicting the intention to interact with them through the device's pick-up detection. Tickle is a wearable interface that can be worn on the user's fingers (as a ring) or fixed to it (with nail polish). Therefore, the device controlled by finger gestures can be any generic object, provided they have an interface for receiving the sensor's signal. The proposed interface is an example towards the idea of ubiquitous computing and the vision of seamless interactions with grasped objects. Finally, she developed an interactive desktop lamp. The perception and gestural output space were explored in two studies those results will be summarized as the "grammar" of the lamp.

Participants

- Tom Erickson
Minneapolis, US
- Elizabeth Gerber
Northwestern University – Evanston, US
- Steven Greenspan
CA Labs, US
- Stefan Greuter
Deakin University – Melbourne, AU
- Jonathan Grudin
Microsoft Research – Redmond, US
- Ti Hoang
RMIT University – Melbourne, AU
- Masahiko Inami
University of Tokyo, JP
- Wendy Ju
Cornell Tech – New York, US
- Kai Kunze
Keio University – Yokohama, JP
- Joseph La Delfa
RMIT University – Melbourne, AU
- Zhuying Li
RMIT University – Melbourne, AU
- Pedro Lopes
Hasso-Plattner-Institut – Potsdam, DE
- Pattie Maes
MIT – Cambridge, US
- Joseph Marshall
University of Nottingham, GB
- Jochen Meyer
OFFIS – Oldenburg, DE
- Florian Mueller
RMIT University – Melbourne, AU
- Suranga Nanayakkara
University of Auckland, NZ
- Jun Nishida
University of Tsukuba, JP
- Marianna Obrist
University of Sussex – Brighton, GB
- Harald Reiterer
Universität Konstanz, DE
- Thecla Schiphorst
Simon Fraser University – Surrey, CA
- Caitlyn E. Seim
Georgia Institute of Technology – Atlanta, US
- Jürgen Steimle
Universität des Saarlandes, DE
- Paul Strohmeier
University of Copenhagen, DK
- Dag Svanaes
NTNU – Trondheim, NO
- Dakuo Wang
IBM T.J. Watson Research Center – Yorktown Heights, US
- Martin Weigel
Honda Research Europe – Offenbach, DE
- Katrin Wolf
HAW – Hamburg, DE

Report from Dagstuhl Seminar 18331

Algorithmic Foundations of Programmable Matter

Edited by

Spring Berman[1], Sándor P. Fekete[2], Matthew J. Patitz[3], and Christian Scheideler[4]

1 **Arizona State University – Tempe, US,** spring.berman@asu.edu
2 **TU Braunschweig, DE,** s.fekete@tu-bs.de
3 **University of Arkansas – Fayetteville, US,** patitz@uark.edu
4 **Universität Paderborn, DE,** scheideler@upb.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 18331, "Algorithmic Foundations of Programmable Matter", a new and emerging field that combines theoretical work on algorithms with a wide spectrum of practical applications that reach all the way from small-scale embedded systems to cyber-physical structures at nano-scale.

The aim of this seminar was to bring together researchers from computational geometry, distributed computing, DNA computing, and swarm robotics who have worked on programmable matter to inform one another about the newest developments in each area and to discuss future models, approaches, and directions for new research. Similar to the first Dagstuhl seminar on programmable matter (16271), we did focus on some basic problems, but also considered new problems that were now within reach to be studied. During this seminar, we were able to achieve a previously unmatched level of intensity of collaboration, in part due to using a new electronic and interactive web-based platform. This has also allowed for continued research among the attendees based on the work begun during the seminar.

Seminar August 12–17, 2018 – http://www.dagstuhl.de/18331
2012 ACM Subject Classification Theory of computation → Computational geometry, Theory of computation → Self-organization, Computer systems organization → Robotics, Computer systems organization → Self-organizing autonomic computing
Keywords and phrases computational geometry, distributed algorithms, DNA computing, programmable matter, swarm robotics
Digital Object Identifier 10.4230/DagRep.8.8.48
Edited in cooperation with Arne Schmidt

1 Executive Summary

Spring Berman
Sándor P. Fekete
Matt Patitz
Christian Scheideler

The term "programmable matter" refers to any substance that can change its physical properties (shape, density, moduli, conductivity, optical properties, etc.) in a programmable fashion. The role of *algorithmic foundations* of programmable matter continues to grow in

importance due to ongoing progress in a wide range of applications. Examples of cutting-edge application areas with a strong algorithmic flavor include *self-assembling systems*, in which chemical and biological substances such as DNA are designed to form predetermined shapes or carry out massively parallel computations; and *swarm robotics*, in which complex tasks are achieved through the local interactions of robots with highly limited individual capabilities, including micro- and nano-robots. Progress in these application areas has been achieved through close collaboration with algorithmic theoreticians, enabling the investigation of fundamental problems related to system geometry using methods from the field of *computational geometry*, and yielding techniques for decentralized computation from the field of *distributed computing*.

A previous Dagstuhl seminar (16271, Algorithmic Foundations of Programmable Matter) had laid the foundations for further progress by bringing together experts from different fields and focusing on expert surveys and breakout groups. We built on the success of that seminar by expanding its focus on particular challenges that arise from the application areas of programmable matter. For this purpose, we brought together a combination of established experts from DNA computing, swarm robotics, computational geometry, and distributed computing. On the senior level, particants included a number of leading authorities who are established in more than one of the mentioned topics; on the junior level, we had a good selection of highly talented scientists who are able to advance the field by specific contributions.

The seminar started with a plenary introduction of all participants, their research areas and their specific challenges and expectations for the seminar. This was followed by a number of plenary sessions, in which experts gave overviews of broad developments and specific open problems.

- Erik Demaine gave an overview of challenges for **geometric algorithms** in the settings of reconfigurable robots (both modular and folding robots that can become any possible shape), robot swarms (which may be so small and simple that they have no identity), and self-assembly (building computers and replicators out of DNA tiles).
- Dave Doty and Chris Thachuk gave a survey of the basics of experimental and theoretical **DNA tile self-assembly**, concluding with suggestions for theoretical problems related to programmable control of the nucleation of assemblies. A second part consisted of a survey of DNA strand displacement, including the problem of orienting molecules on a surface with the use of DNA origami and some clever shapes that can "align" themselves into target placements.
- Andréa Richa presented an overview of **self-organizing particle systems**, describing programmable matter as an abstract collection of simple computational elements (particles) with limited memory that each execute fully distributed, local, asynchronous algorithms to self-organize and solve system-wide problems such as movement, (re)configuration, and coordination.
- Aaron Becker discussed the connection between **robot swarms and programmable matter**, in particular in a setting with a global input to a whole particle swarm, as well as open questions arising from the use of mobile robots to fold 2D planar stock into 3D bricks and to connect the bricks together.

Spread throughout the week, further presentations were given by Spring Berman (applications and open challenges in swarm robotics and a control-theoretic framework for robotic swarms and programmable matter), Julien Bourgeois (realizing programmable matter with modular robots), Luca Cardelli (sequenceable DNA algorithms), Kenneth Cheung

(programmable modular periodic metamaterials), Sándor Fekete (coordinated motion planning), Roderich Groß (capabilities of individual units in distributed robotic systems and making programmable matter self-propel efficiently), Dan Halperin (hard vs. easy tasks in multi-robot motion planning), Heiko Hamann (self-assembly and collective construction based on minimal surprise), Lila Kari (DNA smart-tile self-assembly and computational CRISPR), MinJun Kim (engineering particles for robot swarms and modular microrobotics), Alcherio Martinoli (fluid-mediated stochastic self-assembly), Friedhelm Meyer auf der Heide (continuous strategies for swarm robotics), Nils Napp (autonomous construction in unstructured environments), Pekka Orponen (algorithmic design of RNA nanostructures) and Christian Scheideler (a survey on hybrid programmable matter).

A key feature of the seminar was exceptionally intensive, interdisciplinary collaboration throughout the week, based on the use of the new interactive electronic tool `coauthor`. This tool[1], specifically developed for use in a workshop-like environment, is an excellent platform that provides a versatile medium for collaborative research discussions, and maintains easily accessible structured records for future reference. We have found that `coauthor` greatly facilitated the work done during the seminar, enabling not just identification of, but also dynamic research work on a number of new topics. These include (A) specific problems in the context of hybrid models for programmable matter, in which there is a set of active micro-robots that can move a large set of simple material tiles that cannot move themselves; (B) aspects of distributed boundary detection for self-organizing swarms; (C) fundamental issues related to the computational equivalence of completely different self-assembly systems and robotic models; and (D) questions of self-aligning geometric shapes that would allow more robust methods for DNA origami and self-assembly. For some aspects, we were able to resolve long-standing open problems; for others, we made significant progress that will undoubtedly lead to future publications. As a consequence, the seminar has triggered a number of new collaborations and a variety of followup projects that will undoubtedly contribute to further collaborative research activities.

[1] https://github.com/edemaine/coauthor/

2 Contents

3 Overview of Talks

3.1 Robot Swarms and Programmable Matter

Aaron Becker (University of Houston, US)

Joint work of Manzoor, Sheryl; Kim, Min Jun; Schmidt, Arne; Huang, Li; Krupke, Dominic; Fekete, Sándor; Ike, Rhema; Kantari, Saleh
Main reference Arne Schmidt, Sheryl Manzoor, Li Huang, Aaron T. Becker, Sándor P. Fekete: "Efficient Parallel Self-Assembly Under Uniform Control Inputs", CoRR, Vol. abs/1807.01584, 2018.
URL http://arxiv.org/abs/1807.01584

1. Global inputs, where each agent receives the same control input, are often used for tiny robots because it is difficult to fit power, actuation, and computation in tiny robots. This talk gives an overview of global inputs for steering planar robots to perform manipulation, generate formations, assemble into desired shapes, and build maps.

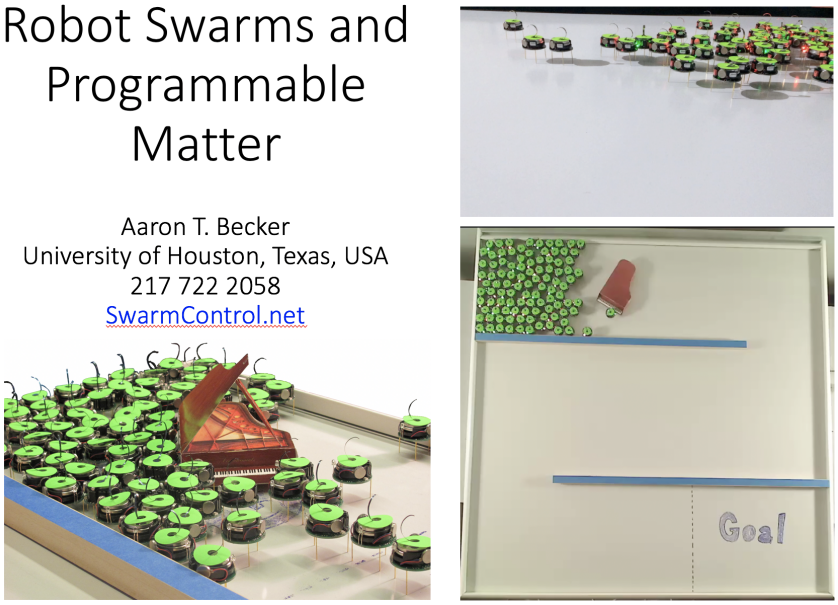

2. It is easier to equip larger robots with computation, sensing, and actuation. The second half of the talk explored open questions using mobile robots to fold 2D planar stock into 3D bricks and connect the bricks together. In this case the planar stock (paper) can be programmed by cutting it, scoring fold lines, and attaching fiducials and handles. Open questions include:
   a.) What classes of shapes can [not] be folded using forces applied to perimeter?
   b.) How can impossible shapes be approximated?
   c.) Minimum number of pushing actions?
   d.) Minimum number of robots?
   e.) How much harder is the distributed version of this problem than the centralized version?

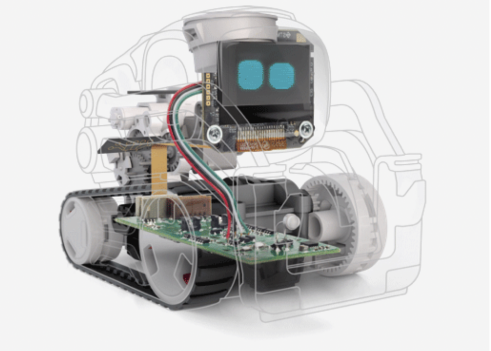

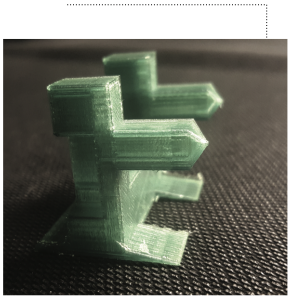

References

1 S. Manzoor, S. Sheckman, J.t Lonsford, H. Kim, Min Jun Kim, and A. Becker, "Parallel Self-Assembly of Polyominoes Under Uniform Control Inputs", IEEE RA-L, 2017, 10.1109/LRA.2017.2715402

2 A. Schmidt, S. Manzoor, L. Huang, A. Becker, S. Fekete, "Efficient Parallel Self- Assembly Under Uniform Control Inputs", IEEE RA-L, 2018. https://arxiv.org/abs/1807.01584

3 P. Keldenich, S. Manzoor, L. Huang, D. Krupke, A. Schmidt, S. Fekete, A. Becker, "On Designing 2D Discrete Workspaces to Sort or Classify Polyominoes", 2018 IEEE IROS 2018

4 A. Mahadev, D. Krupke, S. Fekete, A. Becker, "Mapping and Coverage with a Particle Swarm Controlled by Uniform Inputs", IROS, 2017 10.1109/IROS.2017.8202280

5 Saleh Kantari, Rhema Ike, Aaron Becker, "Designing for Digital Assembly with a Construction Team of Mobile Robots", Association for Computer Aided Design in Architecture (ACADIA) October 18 – 20, 2018, Universidad Iberoamericana, Mexico City

3.2 Swarm Robotics: Applications, Open Challenges, and a Control-Theoretic Framework for Programmable Matter

Spring Berman (Arizona State University – Tempe, US)

Joint work of Spring Berman, Karthik Elamvazhuthi, Andrea L. Bertozzi, Hendrik Kuiper, Matthias Kawski, Fangbo Zhang, Matt Haberland, Ragesh K. Ramachandran, Vaibhav Deshmukh, Shiba Biswal, Zahi Kakish, Chase Adams, Sean Wilson, Theodore P. Pavlic, Ganesh P. Kumar, Aurélie Buffin

Robotic swarms are currently being developed to perform a variety of tasks over large spatial and temporal scales. However, significant technical challenges remain before these systems can be robustly deployed in unstructured, dynamic environments. We are addressing the problem of controlling swarms in scenarios where the robots lack global localization, prior data about the environment, and reliable inter-robot communication. As in natural swarms, the highly resource-constrained robots would be restricted to information obtained through

local sensing and signaling. We are developing a rigorous control and estimation framework, which may be useful for designing programmable matter, for swarms that are subject to these constraints. This framework will enable swarms to operate largely autonomously, with user input consisting only of high-level directives that map to a small set of robot parameters. We use stochastic and deterministic models from chemical kinetics and fluid dynamics to describe the robots' roles, task transitions, and spatiotemporal distributions at both the microscopic (individual) and macroscopic (population) levels. We have applied this framework to design stochastic strategies for coverage, assembly, task allocation, mapping, and scalar field estimation, as well as decentralized, ant-inspired approaches to cooperative manipulation.

References

1 Sean Wilson, Theodore P. Pavlic, Ganesh P. Kumar, Aurélie Buffin, Stephen Pratt, and Spring Berman. "Design of Ant-Inspired Stochastic Control Policies for Collective Transport by Robotic Swarms." Swarm Intelligence, 8(4):303-327, Dec. 2014. DOI: 10.1007/s11721-014-0100-8

2 Karthik Elamvazhuthi, Shiba Biswal, and Spring Berman. "Mean-Field Stabilization of Robotic Swarms to Probability Distributions with Disconnected Supports." American Control Conference (ACC), Milwaukee, WI, 2018. DOI: 10.23919/ACC.2018.8431780

3.3 Realizing Programmable Matter with Modular Robots

Julien Bourgeois (FEMTO-ST Institute – Montbéliard, FR)

Main reference Julien Bourgeois, Benoît Piranda, André Naz, Nicolas Boillot, Hakim Mabed, Dominique Dhoutaut, Thadeu Tucci, Hicham Lakhlef: "Programmable matter as a cyber-physical conjugation", in Proc. of the 2016 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2016, Budapest, Hungary, October 9-12, 2016, pp. 2942–2947, IEEE, 2016.
URL http://dx.doi.org/10.1109/SMC.2016.7844687

Technological advances, especially in the miniaturization of robotic devices foreshadow the emergence of large-scale ensembles of small-size resource-constrained robots that distributively cooperate to achieve complex tasks. These ensembles are formed by independent, intelligent and communicating units which act as a whole ensemble which can be used to build programmable matter, i.e., matter able to change its shape. In my talk, I present our research effort in building Programmable Matter (PM) based on modular robots. To do this, we use micro-technology to scale down the size of each element, and we study geometry, structure, actuation, power, electronics and integration. To manage the complexity of this kind of environment, we propose a complete environment including programmable hardware, a programming language, a compiler, a simulator, a debugger and distributed algorithms.

3.4 Sequenceable DNA Algorithms

Luca Cardelli (Microsoft Research UK – Cambridge, GB)

We give an introduction to DNA Strand Displacement: a technique that is used to program interactions between DNA strands in such a way, e.g., as to emulate the kinetics of an arbitrary finite network of chemical reactions. We discuss current capabilities and trends in DNA nanotechnology, including "high throughput" equipment that can read and write DNA massively in parallel.

High throughput DNA synthesis and sequencing render easily feasible a new class of algorithms that use $O(n^2)$ structures in input and output. We give two examples of such algorithms, for detecting the coincidence of events, and for detecting the preorder of events, over the course of an experiment in a biochemical soup.

3.5 Programmable Modular Periodic Metamaterials

Kenneth C. Cheung (NASA – Moffett Field, US)

Current practice in material and manufacturing efficiency is far from theorized physical limits, as demanded by long term and large scale space exploration applications. Significant progress may be achievable with reversible material assembly. Various programmable matter techniques utilize many of a small set of discrete modules that can be used to compute or program any area or volumetric shape, with any internal pattern. Metamaterials appear to be a promising demonstration target for discretized programmable material methods. In general these materials provide the capability of prescribing unique combinations of material properties to suit custom applications. Mechanical metamaterials are an example of periodic geometry governed materials that have seen near theoretically ideal properties when implemented through modular construction. These materials offer hierarchical decomposition in modeling, with bulk properties that can be predicted from component measurements, and programmed by relative placement of discrete part types with differing properties. Nano- and micro- scale mechanical metamaterials and periodic structures in general are clear targets for molecular assembly (with foundational work already in place). Macro-scale implementation may be a very interesting application for distributed robotics, as physical realizations of discrete theoretical models. These robots have characteristic dimension on the order of that of the modular cell or voxel, and use the regularity of the built assembly to simplify path planning, locomotion and manipulation (with low precision requirements), and allow low numbers of states and degrees of freedom (DOF) per robot. Many interesting open problems exist for determining optimality and complexity of planning and scheduling with various structural geometries and multi-robot system architectures.

References

1 Cheung, K. C., Demaine, E. D., Bachrach, J. R., Griffith, S. (2011). Programmable assembly with universally foldable strings (moteins). IEEE Transactions on Robotics, 27(4), 718–729.
2 Cheung, K. C., Gershenfeld, N. (2013). Reversibly assembled cellular composite materials. Science, 1240889.
3 Jenett, B., Calisch, S., Cellucci, D., Cramer, N., Gershenfeld, N., Swei, S., Cheung, K. C. (2017). Digital morphing wing: active wing shaping concept using composite lattice-based cellular structures. Soft Robotics, 4(1), 33–48.
4 Jenett, B., Cheung, K. (2017). BILL-E: Robotic Platform for Locomotion and Manipulation of Lightweight Space Structures. In 25th AIAA/AHS Adaptive Structures Conference (p. 1876).
5 Jenett, B., Cellucci, D. (2017, May). A mobile robot for locomotion through a 3D periodic lattice environment. In Robotics and Automation (ICRA), 2017 IEEE International Conference on (pp. 5474–5479). IEEE.
6 Gregg, C. E., Kim, J. H., Cheung, K. C. (2018). Ultra-Light and Scalable Composite Lattice Materials. Advanced Engineering Materials.

3.6 Replicators, Transformers, and Robot Swarms: Science Fiction through Geometric Algorithms

Erik D. Demaine (MIT – Cambridge, US)

Science fiction is a great inspiration for science. How can we build reconfigurable robots like Transformers or Terminator 2? How can we build Star Trek-style replicators that duplicate or mass-produce a given shape at the nano scale? How can we orchestrate the motion of a large swarm of robots? Recently we've been exploring possible answers to these questions through computational geometry, in the settings of reconfigurable robots (both modular and folding robots that can become any possible shape), robot swarms (which may be so small and simple that they have no identity), and self-assembly (building computers and replicators out of DNA tiles).

3.7 DNA tile self-assembly

David Doty (University of California – Davis, US)

I survey the basics of experimental and theoretical DNA tile self-assembly, starting with experimental efforts initiate by Ned Seeman to create unbounded period crystal lattices with DNA tiles, experimental efforts initiated by Erik Winfree to create algorithmic self-assembling tiles, and basic theoretical results about the computational ability of the tiles. I conclude with suggestions for theoretical problems on programmable control of nucleation that, if resolved, could be immediately tested in a wet lab, with luck greatly increasing the yield of tile-based self-assembled structures, compared to the current low yield state-of-the-art.

3.8 Coordinated motion planning: Reconfiguring a swarm of labeled robots with bounded stretch

Sándor P. Fekete (TU Braunschweig, DE)

Joint work of Aaron T. Becker, Erik D. Demaine, Sándor P. Fekete, Phillip Keldenich, Matthias Konitzny, Lillian Lin, Henk Meijer, Christian Scheffer

We motivate, visualize and demonstrate recent work for minimizing the total execution time of a coordinated, *parallel* motion plan for a swarm of N robots in the absence of obstacles. Under relatively mild assumptions on the separability of robots, the algorithm achieves *constant stretch*: If all robots want to move at most d units from their respective starting positions, then the total duration of the overall schedule (and hence the distance traveled by each robot) is $O(d)$ steps; this implies constant-factor approximation for the optimization problem. Also mentioned is an NP-hardness result for finding an optimal schedule, even in the case in which robot positions are restricted to a regular grid. On the other hand, we show that for densely packed disks that cannot be well separated, a stretch factor $\Omega(N^{1/4})$ is required in the worst case; we establish an achievable stretch factor of $O(N^{1/2})$ even in this case. We also sketch geometric difficulties of computing optimal trajectories, even for just two unit disks.

References

1 Aaron T. Becker, Sándor P. Fekete, Phillip Keldenich, Matthias Konitzny, Lillian Lin, Christian Scheffer: Coordinated Motion Planning: The Video (Multimedia Exposition). Symposium on Computational Geometry 2018: 74:1–74:6
2 Erik D. Demaine, Sándor P. Fekete, Phillip Keldenich, Christian Scheffer, Henk Meijer: Coordinated Motion Planning: Reconfiguring a Swarm of Labeled Robots with Bounded Stretch. Symposium on Computational Geometry 2018: 29:1–29:15

3.9 Less is more? Defining your building blocks

Roderich Groß (University of Sheffield, GB)

Joint work of Roderich Groß, Melvin Gauci, Jianing Chen, Matthew Doyle, Anil Özdemir

When designing a distributed robotic system to exhibit collective behavior, choices need to be made regarding the capabilities of the underlying individual units. These choices impact on the potential to scale the units down in size and up in numbers. This talk first shows how to design behavioral rules of extreme simplicity. We look at (i) rules emulating granular material, and (ii) rules that require no arithmetic computation. Second, we address the open challenge to make programmable matter self-propel efficiently, by introducing a novel propulsion concept. This leads to robots of arbitrary morphology, which provably move towards a goal in their environment even though they lack arithmetic computation.

References

1 Gauci M., Chen J., Li W., Dodd T. J., Groß R. Self-Organized Aggregation without Computation The International Journal of Robotics Research, 33(8):1145–1161, 2014 http://dx.doi.org/10.1177/0278364914525244
2 Chen J., Gauci M., Li W., Kolling A., and Groß R. Occlusion-based cooperative transport with a swarm of miniature mobile robots IEEE Transactions on Robotics, 31(2):307–321, 2015 http://dx.doi.org/10.1109/TRO.2015.2400731
3 Doyle M., Xu X., Gu Y., Perez-Diaz F., Parrot C., Groß R. Modular hydraulic propulsion: a robot that moves by routing fluid through itself ICRA 2016, IEEE (2016) 5189–5196 http://dx.doi.org/10.1109/ICRA.2016.7487725

3.10 Hard vs. Easy in Multi-Robot Motion Planning

Dan Halperin (Tel Aviv University, IL)

Early results in robot motion planning had forecast a bleak future for the field by showing that problems with many degrees of freedom, and in particular those involving fleets of robots, are intractable. Then came sampling-based planners, which have been successfully, and often easily, solving a large variety of problems with many degrees of freedom. We strive to formally determine what makes a motion-planning problem with many degrees of freedom easy or hard. In the first part of the talk I'll describe our quest to resolve this (still wide open) problem, and some progress we have made in the context of multi-robot motion planning. In the second part of the talk I'll review recent algorithms that we have developed for multi-robot motion planning, which come with near- or asymptotic-optimality guarantees.

References

1 Aviv Adler, Mark de Berg, Dan Halperin, Kiril Solovey: Efficient Multi-Robot Motion Planning for Unlabeled Discs in Simple Polygons. IEEE Trans. Automation Science and Engineering 12(4): 1309–1317 (2015)
2 Andrew Dobson, Kiril Solovey, Rahul Shome, Dan Halperin, Kostas E. Bekris: Scalable asymptotically-optimal multi-robot motion planning. MRS 2017: 120–127

3.11 Self-assembly and collective construction based on minimal surprise

Heiko Hamann (Universität Lübeck, DE)

Joint work of Heiko Hamann, Tanja Kaiser, Richard Borkowski
Main reference Heiko Hamann: "Evolution of Collective Behaviors by Minimizing Surprise", Alife, MIT Press, pp. 344–351, 2014.
URL https://doi.org/10.7551/978-0-262-32621-6-ch055

What about we don't tell our robots what they have to do? With the approach of minimizing surprise, we only ask the robots to create situations that allow for easy predictions. Each robot has a controller (artificial neural network, ANN) to select the next action and a prediction machine (ANN) to predict future sensor input (of the next time step). Here, these pairs of networks are tuned by an evolutionary algorithm. The fitness (reward) is based

on correct predictions only. There is no desired or predefined behavior. Applied to robot swarms, we observe typical swarm behaviors, such as aggregation, dispersion, and flocking. In a self-assembly setup we observe also aggregation, dispersion, line formations, and the formation of triangular lattices.

3.12 (DNA) Smart-tile Self-Assembly and Computational CRISPR

Lila Kari (University of Waterloo, CA)

Joint work of Lila Kari, Amirhossein Simjour
Main reference Lila Kari, Amirhossein Simjour: "Smart Tile Self-Assembly and Replication", Fundam. Inform., Vol. 154(1-4), pp. 239–260, 2017.
URL http://dx.doi.org/10.3233/FI-2017-1564

We present the concept of (in-vitro DNA) self-assembly of smart tiles, i.e., tiles which possess a local computational device, in addition to having edge glues that can be activated or deactivated by signals (joint work with Amir Simjour). The local tile computational device can range from nonexistent, to being a counter, a simple look-up table, a finite state machine, all the way to being a Turing machine. Thus, this model may offer a general framework to discuss and compare various tile self-assembly systems. We demonstrate the potential of self-assembly with smart tiles to efficiently perform robotic tasks such as the replication of convex shapes. The smart tile assembly system that we propose for convex shape replication does not make any assumption on the glues and signals of the interior tiles of the input supertile, adds tiles one at a time, and uses a scaffold to assemble a replica adjacent to the input supertile.

In the second part of the talk, we draw a parallel between the mechanism of targeted gene editing by CRISPR-CAS9 and Contextual Insertion/Deletion Systems, a DNA-inspired computational model with Turing machine computational power. We suggest that a version of contextual ins/del systems could be defined that accurately utilizes the CRISP/CAS9 mechanism, making the latter a suitable candidate for easy, efficient, and customizable in-vivo DNA computation.

3.13 Engineering Particles for Robot Swarms and Modular Microrobotics

MinJun Kim (SMU – Dallas, US)

The realization of reconfigurable modular microrobots could aid drug delivery and microsurgery by allowing a single system to navigate diverse environments and perform multiple tasks. So far, microrobotic systems are limited by insufficient versatility; for instance, helical shapes commonly used for magnetic swimmers cannot effectively assemble and disassemble into different size and shapes. Here by using microswimmers with simple geometries constructed of spherical particles, we show how magnetohydrodynamics can be used to assemble and disassemble modular microrobots with different physical characteristics. We develop a mechanistic physical model that we use to improve assembly strategies. Furthermore, we

experimentally demonstrate the feasibility of dynamically changing the physical properties of microswimmers through assembly and disassembly in a controlled fluidic environment. Finally, we show that different configurations have different swimming properties by examining swimming speed dependence on configuration size.

References

1 U. K. Cheang, F. Meshkati, H. Kim, K. Lee, H. C. Fu, Min Jun Kim, "Versatile microrobotics using simple modular subunits," Sci. Rep., Vol. 6, 30472, 2016.
2 J. Ali, U. K. Cheang, Y. Liu, H. Kim, L. Rogowski, S. Sheckman, P. Patel, W. Sun, Min Jun Kim, "Fabrication and magnetic control of alginate cellular microrobots," AIP Advances, Vol. 6, 125205, 2016.
3 J. Ali, U. K. Cheang, A. Darvish, Min Jun Kim, "Biotemplated flagellar nanoswimmers," APL Materials, Vol. 5, 116106, 2017.
4 H. Kim, Min Jun Kim, "Electric field control of bacteria-powered microrobots (BPMs) using static obstacle avoidance algorithm," IEEE Trans. Robot., Vol. 32, No. 1, p. 125-137, 2016.
5 P. S. S. Kim, A. Becker, Y. Ou, A. A. Julius, Min Jun Kim, "Imparting magnetic dipole heterogeneity to internalized iron oxide nanoparticles for microorganism swarm control," J. Nanopart. Res., Vol. 17, 144, 2015.
6 U. K. Cheang, Min Jun Kim, "Self-assembly of robotic micro- and nanoswimmers using magnetic nanoparticles," J. Nanopart. Res., Vol. 17, 145, 2015.
7 O. Yan, D. H. Kim, P. S. S. Kim, Min Jun Kim, A. A. Julius, "Motion control of magnetized Tetrahymena pyriformis cells by magnetic field with model predictive control (MPC)," Int. J. Robotics Res, Vol. 32, No. 1, p129-139, 2013.
8 E. B. Steager, M. S. Sakar, D. H. Kim, V. Kumar, G. Pappas, Min Jun Kim, "Electrokinetic and optical control of bacterial microrobots," J. Micromech. Microeng., Vol. 21, 035001, 2011.

3.14 Fluid-Mediated Stochastic Self-Assembly: Towards Bridging Centimetric and Submillimetric Scales

Alcherio Martinoli (EPFL – Lausanne, CH)

Joint work of Bahar Haghighat, Massimo Mastrangeli, Grégory Mermoud, Felix Schill, Alcherio Martinoli
Main reference Bahar Haghighat, Massimo Mastrangeli, Grégory Mermoud, Felix Schill, Alcherio Martinoli: "Fluid-Mediated Stochastic Self-Assembly at Centimetric and Sub-Millimetric Scales: Design, Modeling, and Control", Micromachines, Vol. 7(8), p. 138, 2016.
URL http://dx.doi.org/10.3390/mi7080138

Miniature robots at centimeter scale can be effective demonstrators: they can be designed and manufactured leveraging off-the-shelf components and standard mechatronic recipes. Unfortunately, the application areas for this scale are limited while many potential applications are available at the submillimeter scale. Devices at the submillimeter scale cannot be endowed with similar resources as their centimetric counterpart as the manufacturing technology at this scale is still very expensive, needs hardware customization, and multiple functionalities are difficult to integrate in a single device with canonical top-down manufacturing techniques. One of the promising techniques to manufacture more complex microrobots is to leverage existing Micro-Electro-Mechanical Systems (MEMS) achieving different functionalities, produced with standard micromachining procedures, and use them as building blocks for a fluid-mediated

self-assembling process. In order to efficiently guide the self-assembly process, stochastic modeling and control techniques developed for centimeter devices can help. In particular, in this talk I illustrate a vision-based closed-loop framework able to both automatically create models at multiple abstraction levels and optimize the agitation of the liquid in which the self-assembly process takes place for both passive centimeter and sub-millimeter building blocks. I then describe also recent self-assembly results achieved with a demonstrator having the same dimensions of the previous one at centimeter scale but consisting of robotic modules endowed with a programmable ruleset and able to control their on-board latching properties. While these latter properties are difficult to transpose to submillimeter modules, this successive research effort has allowed us to develop an effective software framework for self-assembly experiments and gain further insight in distributed control strategies potentially deployable at submillimeter level in the future.

References

1 Bahar Haghighat and Alcherio Martinoli *Automatic synthesis of rulesets for programmable stochastic self-assembly of rotationally symmetric robotic modules*. Swarm Intelligence Journal, 11(3-4): 243-270, 2017
2 Bahar Haghighat, Massimo Mastrangeli, Gregory Mermoud, Felix Schill, and Alcherio Martinoli *Fluid-Mediated Stochastic Self-Assembly at Centimetric and Sub-Millimetric Scales: Design, Modeling, and Control*. Micromachines, 7(8): 138, 2016
3 Massimo Mastrangeli, Felix Schill, Jonas Goldowsky, Helmut Knapp, Juergen Brugger, and Alcherio Martinoli *Automated Real-Time Control of Fluidic Self-Assembly of Microparticles*. Proc. of the 2014 IEEE Int. Conf. on Robotics and Automation, May-June 2014, Hong Kong, China, pp. 5860-5865
4 Gregory Mermoud, Massimo Mastrangeli, Utkarsh Upadhyay, and Alcherio Martinoli *Real-Time Automated Modeling and Control of Self-Assembling Systems*. Proc. of the 2012 IEEE Int. Conf. on Robotics and Automation, May 2012, Saint Paul, MN, USA, pp. 4266-4273

3.15 Continuous Strategies for Swarm Robotics

Friedhelm Meyer auf der Heide (Universität Paderborn, DE)

Joint work of Friedhelm Meyer auf der Heide, Peter Kling, Barbara Kempkes, Pavel Podlipyan and many others.
Main reference Bastian Degener, Barbara Kempkes, Peter Kling, Friedhelm Meyer auf der Heide: "Linear and Competitive Strategies for Continuous Robot Formation Problems", TOPC, Vol. 2(1), pp. 2:1–2:18, 2015.
URL http://dx.doi.org/10.1145/2742341

I consider large swarms of relatively simple mobile robots deployed to the plane. Each robot has only very limited local information about the swarm. More precisely, a swarm consists of identical, anonymous robots: they are points in the plane and their local information consists only of the relative positions of the robots within a small, bounded viewing radius. My focus is on strategies of such swarms that result in formations problems like "gathering in one point". I present several strategies for such formation problems, and discuss upper and lower bounds for their runtime.

First, I introduce continuous strategies, where each robot continuously observes its neighborhood and continuously adapts its speed (with given speed limit) and direction following a local rule. I present the class of contracting strategies, show that they perform

gathering in quadratic time, and present best and worst case examples. With the go-on-bisector-strategy, I present a time optimal gathering strategy. Finally, I introduce the variant of the go-to-the center strategy that only considers edges of the Gabriel subgraph of the unit-disk graph, and give experimental evidence that it almost never produces early collisions. Most of the presented work is based on the following two publications:

References

1 Shouwei Li, Christine Markarian, Friedhelm Meyer auf der Heide, Pavel Podlipyan: A Continuous Strategy for Collisionless Gathering. ALGOSENSORS 2017: 182-197
2 Bastian Degener, Barbara Kempkes, Peter Kling, Friedhelm Meyer auf der Heide: Linear and Competitive Strategies for Continuous Robot Formation Problems. TOPC 2(1): 2:1-2:18 (2015)

3.16 Autonomous Construction in Unstructured Environments

Nils Napp (Buffalo State – The State University of New York, US)

The talk presents research on reliably building support structures in cluttered and un-prepared environments using autonomous robots, both from a theoretical and practical perspective. The focus of the talk is modeling the physical world and the robot's interactions with it, in such a way that allows the system to synthesize correct actions on the fly based on local sensor information. Ideally, autonomous construction systems should be able to use a wide variety of building materials, yet the ability to plan and act is closely tied to particular building materials. The talk presents some of the challenges of modeling different material types and how they affect a robot's ability to formulate long-term plans.

3.17 Algorithmic design of RNA nanostructures

Pekka Orponen (Aalto University, FI)

Joint work of Pekka Orponen, Antti Elonen, Ibuki Kawamata, Lukas Oesinghaus, Abdulmelik Mohammed, Jani Seitsonen, Yuki Suzuki, Friedrich C. Simmel

Inspired by the remarkable success of DNA nanotechnology, using RNA as nanoscale construction material is attracting increasing attention. The promise of RNA is that it can be produced in quantity by natural processes at room temperature, e.g. by polymerase transcription in cell culture. The challenge, on the other hand, is that the folding process of single-stranded RNA is at the moment much less well understood than that of double-stranded DNA, and there are not similarly well-established robust design protocols as the origami technique in DNA nanotechnology.

Following a brief introduction to some basic concepts in RNA nanotechnology, we present an approach to single-stranded RNA self-assembly of general polyhedral shapes. The technique is based on first routing the RNA strand twice around a spanning tree of the polyhedron's mesh skeleton, in order to create stem helices representing the spanning-tree edges, and then complementing the design by building the non-spanning tree edges from kissing loop motifs. A design tool to support this protocol has been implemented and some initial designs synthesised and imaged.

3.18 Self-organizing Particle Systems

Andréa Richa (Arizona State University – Tempe, US)

Joint work of Andréa Richa, Sarah Cannon, Joshua Daymude, Zahra Derakhshandeh, Robert Gmyr, Cem Gokmen, Kristian Hinnenthal, Irina Kostitsyna, Shengkai Li, Alexandra Porter, Dana Randall, Will Savoye, Christian Scheideler, Thim Strothmann
URL https://sops.engineering.asu.edu/

Many programmable matter systems have been developed, including modular and swarm robotics, synthetic biology, DNA tiling, and smart materials. We describe programmable matter as an abstract collection of simple computational elements (particles) with limited memory that each execute fully distributed, local, asynchronous algorithms to self-organize and solve system-wide problems such as movement, configuration, and coordination. Self-organizing particle systems (SOPS) have many interesting applications like coating objects for monitoring and repair purposes, and forming nano-scale devices for surgery and molecular-scale electronic structures. In this talk, we describe our work on establishing an algorithmic foundation for programmable matter. We investigate how macro-scale system behaviors can naturally emerge from local micro-behaviors by individual particles. We start by investigating shape formation, leader election and coating in SOPS. We then utilize tools from statistical physics and Markov chain analysis to translate Markov chains defined at a system level into asynchronous, distributed, local algorithms for self-organizing particle systems that drive the emergent phenomenon of compression, expansion, bridging, separation, and phototaxing, also establishing direct ties to the notion of "active matter" in physics. Ongoing work also addresses the convex hull problem in the context of self-organizing particle systems.

3.19 Survey on Hybrid Programmable Matter

Christian Scheideler (Universität Paderborn, DE)

Joint work of Robert Gmyr, Kristian Hinnenthal, Irina Kostitsyna, Fabian Kuhn, Dorian Rudolph, Thim Strothmann, Christian Scheideler

In my survey I present a new model for programmable matter that consists of simple nano-robots acting on stateless tiles. Computationally, the nano-robots are only as powerful as finite automata. A nano-robot acts in look-compute-move cycles: it first looks at its immediate neighborhood to determine which positions around it are occupied by tiles, and based on that and its current state it decides on which state to switch to and which move to perform. As a move, a nano-robot may either just move to a neighboring position, pick up a tile below it, or place a tile it is carrying at the position below it. I show that even with just a single nano-robot various shape formation and shape detection problems can be solved. The results presented in this talk appear at the MFCS 2018 and DNA 2018 conferences [1, 2].

References

1 Robert Gmyr, Kristian Hinnenthal, Irina Kostitsyna, Fabian Kuhn, Dorian Rudolph, and Christian Scheideler. Shape Recognition by a Finite Automaton Robot. In Proceedings Mathematical Foundations of Computer Science (MFCS 2018), pp. 52:1–52:15.

2 Robert Gmyr, Kristian Hinnenthal, Irina Kostitsyna, Fabian Kuhn, Dorian Rudolph, Christian Scheideler, and Thim Strothmann. Forming Tile Shapes with Simple Robots. In Proceedings DNA Computing and Molecular Programming (DNA 2018), pp. 122–138.

3.20 Using DNA to compute and to organize molecules on a surface

Chris Thachuk (University of California – Davis, US)

Joint work of David Kirkpatrick, Ashwin Gopinath, Paul W. K. Rothemund, Erik Winfree

The promise of molecular programming lies in its ability to not only process information autonomously, but to do so in a biochemical context in order to sense and actuate matter. The most sophisticated molecular computing systems that have been experimentally realized have been built upon the DNA strand displacement (DSD) primitive, where a soup of rationally designed nucleotide sequences interact, react, and recombine over time in order to carry out complex computation. After giving a survey of DSD we present the problem of absolutely orienting molecules on a surface with the use of DNA origami and some clever shapes that can 'align' themselves into a target placement.

Participants

- Aaron Becker
University of Houston, US
- Spring Berman
Arizona State University – Tempe, US
- Julien Bourgeois
FEMTO-ST Institute – Montbéliard, FR
- Luca Cardelli
Microsoft Research UK – Cambridge, GB
- Kenneth C. Cheung
NASA – Moffett Field, US
- Joshua J. Daymude
Arizona State University – Tempe, US
- Erik D. Demaine
MIT – Cambridge, US
- David Doty
University of California – Davis, US
- Sándor Fekete
TU Braunschweig, DE
- Roderich Gross
University of Sheffield, GB
- Dan Halperin
Tel Aviv University, IL
- Heiko Hamann
Universität Lübeck, DE
- Kristian Hinnenthal
Universität Paderborn, DE
- Lila Kari
University of Waterloo, CA
- MinJun Kim
SMU – Dallas, US
- Irina Kostitsyna
TU Eindhoven, NL
- Dominik Krupke
TU Braunschweig, DE
- Alcherio Martinoli
EPFL – Lausanne, CH
- Friedhelm Meyer auf der Heide
Universität Paderborn, DE
- Othon Michail
University of Liverpool, GB
- Joseph S. B. Mitchell
Stony Brook University, US
- Nils Napp
Buffalo State – The State University of New York , US
- Ram Prasadh Narayanan
Ben Gurion University – Beer Sheva, IL
- Pekka Orponen
Aalto University, FI
- Matthew J. Patitz
University of Arkansas – Fayetteville, US
- Andréa Richa
Arizona State University – Tempe, US
- Marcel J. M. Roeloffzen
TU Eindhoven, NL
- Kay Römer
TU Graz, AT
- Trent Rogers
University of Arkansas – Fayetteville, US
- Dorian Rudolph
Universität Paderborn, DE
- Christian Scheideler
Universität Paderborn, DE
- Stefan Schmid
Universität Wien, AT
- Arne Schmidt
TU Braunschweig, DE
- Chris Thachuk
California Institute of Technology – Pasadena, US
- Pierre Thalamy
FEMTO-ST Institute – Montbéliard, FR
- André van Renssen
The University of Sydney, AU
- Jennifer L. Welch
Texas A&M University – College Station, US

Report from Dagstuhl Seminar 18332

Blockchain Technology for Collaborative Information Systems

Edited by

Marlon Dumas[1], Richard Hull[2], Jan Mendling[3], and Ingo Weber[4]

1 **University of Tartu, EE,** marlon.dumas@ut.ee
2 **IBM TJ Watson Research Center – Yorktown Heights, US,** hull@us.ibm.com
3 **Wirtschaftsuniversität Wien, AT,** jan.mendling@wu.ac.at
4 **Data61, CSIRO – Sydney, AU,** ingo.weber@data61.csiro.au

Abstract

Blockchain technology enables an evolving set of parties to maintain a safe, permanent, and tamper-proof ledger of transactions without a central authority. This technology opens manifold opportunities to redesign business-to-business collaborations, while bringing about numerous challenges. These opportunities and challenges were discussed in the Dagstuhl Seminar 18332 "Blockchain Technology for Collaborative Information Systems". This report documents the program and the outcomes of the seminar.

Seminar August 12–17, 2018 – http://www.dagstuhl.de/18332
2012 ACM Subject Classification Applied computing → Business process management, Information systems → Collaborative and social computing systems and tools
Keywords and phrases Blockchain, BPM, Business Collaboration, Commerce, Logistics, Business Models (economic), Smart Contracts, Privacy
Digital Object Identifier 10.4230/DagRep.8.8.67

1 Executive summary

Marlon Dumas (University of Tartu, EE)
Richard Hull (IBM TJ Watson Research Center – Yorktown Heights, US)
Jan Mendling (Wirtschaftsuniversität Wien, AT)
Ingo Weber (Data61, CSIRO – Sydney, AU)

Blockchain technology enables an evolving set of parties to maintain a safe, permanent, and tamper-proof ledger of transactions without a central authority. This technology opens manifold opportunities to redesign Business-to-Business (B2B) collaborations in a wide range of fields, including supply chain, logistics, service agreements, healthcare, and Industry 4.0. Importantly, it can enable substantial efficiency gains in terms of cost and time it takes to set-up and perform collaborative processes, particularly in settings where there is a lack of trust between the parties involved in the collaboration. Traditionally, collaborative processes are executed by relying on trusted third-party providers such as Electronic Data Interchange (EDI) hubs or escrows. This centralized architecture creates entry barriers and hinders bottom-up innovation. Blockchains and smart contracts enable these processes to be executed in a distributed manner without delegating trust to central authorities nor requiring mutual trust between each pair of parties. Further, blockchain enables fine-grained access

Blockchain Technology for Collaborative Information Systems, *Dagstuhl Reports*, Vol. 8, Issue 08, pp. 67–129
Editors: Marlon Dumas, Richard Hull, Jan Mendling, and Ingo Weber

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

control, thus allowing multiple parties to selectively share their data with each other and to selectively grant permissions to perform transactions on these data.

While blockchain opens up new possibilities, it also raises a number of challenges because it requires us to re-think the way B2B collaborations are designed and implemented. In contrast to centralized collaborative processes, the transparent and decentralized nature of blockchains brings in new challenges related to compliance, control, and privacy, in addition to major scalability and performance challenges. This seminar brought together established and young researchers with forward-thinking industry representatives from both large and start-up companies, in order to establish a research roadmap for blockchain-based collaborative information systems, and to initiate concrete research collaborations between participants along this roadmap.

2 Table of Contents

3 Overview of Talks

This section provides an overview of all the talks held by participants during the seminar.

3.1 Collaborative Business Process Execution on Blockchain: The Caterpillar System

Marlon Dumas (University of Tartu, EE) and Luciano García-Bañuelos (University of Tartu, EE)

Joint work of Orlenys López-Pintado, Luciano García-Bañuelos, Marlon Dumas, Ingo Weber, Alexander Ponomarev
Main reference Orlenys López-Pintado, Luciano García-Bañuelos, Marlon Dumas, Ingo Weber, Alexander Ponomarev: "CATERPILLAR: A Business Process Execution Engine on the Ethereum Blockchain", CoRR, Vol. abs/1808.03517, 2018.
URL https://arxiv.org/abs/1808.03517

Blockchain platforms allow a set of actors to maintain a ledger of transactions without relying on a central authority and to deploy scripts, called smart contracts, that are executed whenever certain transactions occur. These features can be used as building blocks to support the execution of collaborative business processes between mutually untrusting parties. However, implementing business processes using the low-level primitives provided by blockchain platforms is cumbersome and error-prone. In contrast, established business process management systems, such as those based on the standard Business Process Model and Notation (BPMN), provide convenient abstractions for rapid development of process-oriented applications. In this talk, we show how to combine the advantages of a business process management system with those of a blockchain platform. Specifically, we present a BPMN execution engine, namely Caterpillar, designed to support collaborative business processes on top of a blockchain platform. Like any BPMN execution engine, Caterpillar supports the creation of instances of a process model and allows users to monitor the state of process instances and to execute tasks thereof. The specificity of Caterpillar is that the state of each process instance is maintained on the (Ethereum) blockchain and the workflow routing is performed by smart contracts generated by a BPMN-to-Solidity compiler. The Caterpillar compiler supports a large array of BPMN constructs, including subprocesses, multi-instances activities and event handlers.

3.2 Shared Ledger Business Collaboration Language

Rick Hull (IBM TJ Watson Research Center – Yorktown Heights, US)

Blockchain offers the possibility of a fundamentally new way to use information processing in support of business collaboration. We may see a transition from collaboration based on families of binary relationships managed through messaging, to collaboration based on holistic groups of organizations guided by a single "source-of-truth" data repository and shared processing logic. Current approaches to business process and operations management will be extended and transformed as the possibilities and benefits of the single source of

truth are understood and leveraged. This will bring opportunities to re-think current process models and best practices for reducing the models to executables. It also raises challenges in the area of "on-boarding" companies into the usage of Blockchain, because techniques are needed to enable smooth integration between Blockchain-hosted processing and the legacy processing that remains off of Blockchain.

3.3 Assessing the impact of emerging information technologies on processes

Jan Mendling (Wirtschaftsuniversität Wien, AT)

Joint work of Jan Mendling, Gero Decker, Richard Hull, Hajo Reijers, Ingo Weber
Main reference Mendling, Jan and Decker, Gero and Richard, Hull and Hajo A., Reijers and Ingo, Weber: "How do Machine Learning, Robotic Process Automation, and Blockchains Affect the Human Factor in Business Process Management?" Communications of the Association for Information Systems, 43 (Art.19). pp. 297–320, 2018.
URL https://aisel.aisnet.org/cais/vol43/iss1/19/

This talk discusses the impact of emerging technologies on the way how processes can be executed. It makes the point that the often asked naive question "can we do this on the blockchain?" is misleading. Those blockchain variants that come with a Turing-complete programming language can capture any representation of states and transitions.

References

1 Jan Mendling, et al.: Blockchains for business process management-challenges and opportunities. ACM Transactions on Management Information Systems (TMIS) 9.1 (2018): 4.

3.4 Software Architecture and Engineering for Blockchain Applications

Ingo Weber (Data61, CSIRO – Sydney, AU)

Main reference Xiwei Xu, Ingo Weber, and Mark Staples. Architecture for Blockchain Applications. Springer, 2019.

Blockchain technology is emerging and impactful, but its functional and non-functional influence on software applications was initially not well understood. My team and I have been addressing this problem over the past three years, and published a number of papers on various aspects of it. In this talk, I will give an overview over some of our work. As such, I will first present the terminological definitions of blockchain as a concept, technology, network, etc. from our forthcoming book. Then I will provide an overview of the roles blockchain can play in software applications, and what its non-functional properties are. Subsequently I will discuss our design process for choosing a blockchain and configuration. This process starts with the assessment of suitability of blockchain technologies for a given context. Finally, I will provide a short overview of our blockchain design patterns collection. Topics of model-driven engineering and our empirical work are covered in other talks at this seminar, and will therefore not be part of my talk.

3.5 Research for the engineering of blockchain-based systems

Mark Staples (Data61, CSIRO – Sydney, AU)

We will increasingly rely on blockchains (including distributed ledger technologies more broadly) for critical services. Blockchains have the potential to create, and thus put at risk, significant business value. Increasingly blockchains are also being planned to be used for safety-critical applications such as for health and Internet of Things (IoT). So it is important for researchers to not only understand how to design blockchain-based systems, but also to create evidence that they will function correctly, so we can use trustworthy blockchain-based systems.

Functionally, blockchains are a kind of database (an append-only ledger) and a compute platform (executing smart contracts). However, conventional database and compute platforms are owned, operated, and/or administered by single organisations, who become a single point of technical or business failure for the platform. Blockchains are instead operated by a collective. For public blockchains the collective is a large group of anonymous contributors, and in a consortium blockchain the collective is usually defined by contractual arrangements between organisations. Blockchains support efficient and trustworthy ways for organisations to work together. Blockchains also support the representation and control of digital assets, which (unlike normal information assets) can be transferred as forms of exclusive property held by individual parties. Blockchains are potentially disruptive because for the basic services supporting transactional business relationships in industry and society, we can now choose to rely on the neutral territory provided by blockchains, instead of relying on trusted third-parties to facilitate those relationships.

Blockchains have non-functional differences to conventional databases and compute platforms, which impact the architectural design of blockchain-based systems. For example, blockchains have very strong support for integrity, but struggle to directly support confidentiality. Good architectural design is critical to allow systems to benefit from the strengths of blockchain platform, and shore up their weak areas using other off-chain components.

I describe some of the blockchain research from Data61 (CSIRO) that addresses these challenges. This includes the Lorikeet model-driven development platform for generation of blockchain-based systems for business process execution and monitoring, and for control of data-centric business objects in registries. The back-end targets for system generation include Ethereum and Hyperledger Fabric. Our research has also explored the use of these process models for simulation-based performance prediction and for cost analysis. This model-driven generation of systems can take advantage of experience in blockchain architectural design which we have begun to capture in blockchain design patterns. Data61's other blockchain research includes work towards the mechanised formal verification of smart contracts on Ethereum's Virtual Machine, and work towards using declarative representations of legal contracts, with the vision to be used either as specifications for smart contracts, or else as smart contracts, directly interpreted on blockchain platform infrastructure. In other research, we have begun to explore frameworks to integrate support for GS1's EPCIS event data standard into blockchain applications. Although most of our research is on the level of blockchain-based systems, we also have some research activities on the level of underlying blockchain platform, through the development of high-throughput, low-latency consensus mechanisms with a conventional transaction commit semantics, realised in the Red Belly Blockchain. Other platform-level research investigates the use of blockchain principles for IoT networks.

3.6 HyperPubSub: a Decentralized, Permissioned, Publish/Subscribe Service using Blockchains

Kaiwen Zhang (ETS – Montreal, CA), Hans-Arno Jacobsen (TU München, DE), and Nejc Zupan

Joint work of Nejc Zupan, Kaiwen Zhang, Hans-Arno Jacobsen
Main reference Nejc Zupan, Kaiwen Zhang, Hans-Arno Jacobsen: "Hyperpubsub: a decentralized, permissioned, publish/subscribe service using blockchains: demo", in Proc. of the 18th ACM/IFIP/USENIX Middleware Conference: Posters and Demos, Las Vegas, NV, USA, December 11–15, 2017, pp. 15–16, ACM, 2017.
URL https://doi.org/10.1145/3155016.3155018

Since the introduction of Bitcoin in 2008, blockchain systems have evolved immensely in terms of performance and usability. There is a massive focus on building enterprise blockchain solutions, with providers such as IBM and Microsoft already providing Blockchain-as-a Service (BaaS). To facilitate the adoption of blockchain technologies across various business verticals, we argue that middleware plays an integral role in accelerating the development of automated business processes (i.e., smart contracts). We argue that decentralized messaging is a key requirement of many distributed applications and should be provided as a reusable blockchain middleware. Our system, called HyperPubSub, provides decentralized publish/subscribe messaging for a multi-federated, permissioned, environment. HyperPubSub provides secure and privacy-preserving messaging, which is audited using blockchains for validation and monetization purposes. We demonstrate our implementation using Kafka and Hyperledger.

3.7 Tracking Business Processes on the Blockchain

Claudio Di Ciccio (Wirtschaftsuniversität Wien, AT)

Blockchain technology opens up new opportunities for Business Process Management. This is mainly due to its unprecedented capability to let transactions be automatically executed and recorded by Smart Contracts in multi-peer environments, in a decentralised fashion and without central authoritative players to govern the workflow. In this way, blockchains also provide traceability. Traceability of information plays a pivotal role particularly in those supply chains where multiple parties are involved and rigorous criteria must be fulfilled to lead to a successful outcome. In this talk, we investigate how to run a business process in the context of a supply chain on a blockchain infrastructure so as to provide full traceability of its run-time enactment. Our approach retrieves information to track process instances execution solely from the transactions written on-chain. To do so, hash-codes are reverse-engineered based on the Solidity Smart Contracts' encoding of the generating process. We show the results of our investigation by means of an implemented software prototype, with a case study on the reportedly challenging context of the pharmaceutical supply chain.

3.8 Lorikeet: A Model-Driven Engineering Tool for Blockchain-Based Business Process Execution and Asset Management

Qinghua Lu (Data61, CSIRO – Sydney, AU)

Joint work of An Binh Tran, Qinghua Lu, Ingo Weber
Main reference An Binh Tran, Qinghua Lu, Ingo Weber: "Lorikeet: A Model-Driven Engineering Tool for Blockchain-Based Business Process Execution and Asset Management", in Proc. of the Dissertation Award, Demonstration, and Industrial Track at BPM 2018 co-located with 16th International Conference on Business Process Management (BPM 2018), Sydney, Australia, September 9-14, 2018., CEUR Workshop Proceedings, Vol. 2196, pp. 56–60, CEUR-WS.org, 2018.
URL http://ceur-ws.org/Vol-2196/BPM_2018_paper_12.pdf

Blockchain has attracted a broad range of interests including startups, enterprises and governments. A large amount of projects have been conducted to explore how to use blockchain to re-architect systems and to build new applications and business models. However, blockchain is a new technology with still limited tooling and documentation, so there can be a steep learning curve for developers. In addition, it is difficult to fix bugs by releasing new versions of smart contracts and mistakes in smart contracts have led to massive economic loss, such as the DAO exploit. Thus, in this seminar, we demonstrate a model-driven development tool, named Lorikeet, which can automatically generate smart contracts from business process models and/or registry data schema. The architecture of Lorikeet, which consists of a BPMN and registry modeller, smart contract generator and blockchain trigger. The BPMN and registry modeller is presented as a web application for users to build business process and registry models. Business processes are modelled in the Business Process Model and Notation (BPMN) 2.0, while registries are modelled as XML schema. The smart contract generator consists of the BPMN translator and the registry generator. The BPMN translator can automatically create smart contracts in Solidity from BPMN models while the registry generator can produce smart contract based on the registry models. The blockchain trigger communicates with a blockchain node and handles compilation, deployment and interactions with smart contracts. Lorikeet is a well-tested development tool that is used for producing blockchain smart contracts in industry and academia.

3.9 Healthcare Data Management Using Blockchain: Open Challenges and Lessons Learned

Alevtina Dubovitskaya (EPFL – Lausanne, CH)

Joint work of Alevtina Dubovitskaya, Zhigang Xu, Samuel Ryu, Michael Schumacher, Fusheng Wang
Main reference Alevtina Dubovitskaya, Zhigang Xu, Samuel Ryu, Micheal Schumacher, and Fusheng Wang. "Secure and Trustable Electronic Medical Records Sharing using Blockchain." AMIA Annual Symposium Proceedings, Washington, DC, USA, November 4–8, 2017, 2017:650–659, 2018.

Healthcare data are sensitive but have to be frequently shared among the peers in the healthcare network. Blockchain is a distributed ledger technology that may execute arbitrary, programmable transaction logic in the form of smart contracts, and provides a shared, immutable, and transparent append-only register of all the actions that have happened in the network. This provides a unique opportunity to employ blockchain technology to facilitate and enhance healthcare data management for all the actors in the healthcare network. The

immense development of the blockchain technology is happening over the last years, and the interest to the potential of its application in healthcare is growing. However, a quest for such system that guarantees privacy, security, scalability, and efficiency continues. In this talk we discuss potential applications of blockchain in health and present a prototype for blockchain-based consent-management system. We discuss the challenges of adopting such system in medical practice.

3.10 Obsidian: A Safer Blockchain Programming Language

Michael Coblenz (Carnegie Mellon University – Pittsburgh, US)

Blockchain programs have been repeatedly exploited by hackers due to security vulnerabilities, or otherwise found incorrect or unsafe. Verification is one approach toward correct code, but so far, verification is an impractical technique for most programmers. In contrast, Obsidian is a new language that provides particular safety guarantees via an expressive, strong type system. Obsidian allows users to lift state information into types so that the compiler can enforce certain kinds of protocol adherence. Obsidian also supports reasoning about resources, expressing these with linear types.

Obsidian is designed in a user-centered way and serves as a testbed for user-centered language design techniques. We incorporate formative HCI methods to make it more likely that the language is as effective as possible for programmers. We plan to use summative HCI methods in the hope of showing that Obsidian is a more effective tool than prior languages for users to write blockchain programs.

3.11 Ethereum-Based execution of DMN decisions

Stephan Haarman (Hasso-Plattner-Institut – Potsdam, DE)

In business collaborations, participants interact and cooperate to reach business goals. Therefore, their business processes are interconnected via message exchange. Successful collaboration, however, does not only depend on the exchange of messages. Furthermore, local processes must comply with constraints and rules. These are traditionally given by a contract, but conflicts (such as a misunderstanding of the terms) are detected lately, and resolving them becomes expensive. Blockchain technology can help by providing a single source of truth for the data and executable, unambiguous terms (logic).

Using the Decision Model and Notation standard, stakeholders can model business rules precisely. This demo presents a compiler that translates a DMN decision model into smart contract code. The smart contract can be deployed on the Ethereum blockchain: it represents an agreement between participants. An instant specific state (contract) is a single source of truth for all the data. Consequently, all participants can rely on a shared version of data and logic. If all participants comply, then the shared information prevents conflicts. If a conflict occurs, the on-chain data helps to resolve it.

3.12 Collaboration among Adversaries: Distributed Declarative Workflow Execution on a Blockchain

Søren Debois (IT University of Copenhagen, DK)

We study distributed declarative workflow execution in an adversarial setting. In this setting, parties to an agreed-upon workflow do not trust each other to follow that workflow, or suspect the other party might misrepresent proceedings at a later time. We demonstrate how distributed declarative workflow execution can be implemented as smart contracts, guaranteeing (I) enforcement of workflow semantics, and (II) an incontrovertible record of workflow execution history. Crucially, we achieve both properties without relying on a trusted third party. The implementation is based on the Ethereum blockchain, inheriting the security properties (I) and (II) from the guarantees given by that chain. A recurring challenge for both the implementation and the analysis is the cost of operations on Ethereum: This cost must be minimised for honest parties, and an adversary must be prevented from inflicting extra cost on others

3.13 Ergo: A Strongly Typed DSL for Smart Legal Contracts

Jerome Simeon (Clause Inc.– New York, US)

Smart contracts are not legal unless they are legal contracts. Legal contracts are not smart unless they are made computable. Smart legal contracts should not only be legal and smart, but also safe and portable. Ergo is a new DSL which aims to have all those properties. It is developed as part of the Accord Project consortium and open source. It is strongly typed and has a reference implementation and compiler written using the Coq proof assistant.

3.14 Introduction to Hyperledger Fabrics

Petr Novotny (IBM TJ Watson Research Center – Yorktown Heights, US)

Joint work of Angel Nunez Mencias, Donna Dillenberger, D; Petr Novotny, Fabian Toth, Thomas E. Morris
Main reference Angel Nuñez Mencias, Donna N. Dillenberger, Petr Novotny, Fabian Toth, Thomas E. Morris, Volodymyr Paprotski, John C. Dayka, Tamas Visegrady, Bill O'Farrell, Jakob Lang, Ellen Carbarnes: "An optimized blockchain solution for the IBM z14", IBM Journal of Research and Development, Vol. 62(2/3), p. 4, 2018.
URL https://doi.org/10.1147/JRD.2018.2795889

Hyperledger is an open source project hosted by Linux Foundation, supporting collaborative development of blockchain technologies. IBM is the key contributor to the Hyperledger Fabric and to other Hyperledger projects. In this talk, we will explore the key features of permissioned blockchain platform Hyperledger Fabric and of other Hyperledger projects, with the focus on the transaction processing and consensus architecture and protocols, and on the privacy and access control mechanisms. We will then look at a number of use cases which

leverage the permissioned blockchain features in solutions for supply chain, international trade, finance, and others. Finally, we will explore IBM technologies designed for rapid development, simple operation and high performance of blockchain networks.

3.15 Blockchain Research: Process Verification and Beyond

Stefan Schulte (TU Wien, AT)

Joint work of Christoph Prybila, Stefan Schulte, Christoph Hochreiner, Ingo Weber
Main reference Christoph Prybila, Stefan Schulte, Christoph Hochreiner, and Ingo Weber: "Runtime Verification for Business Processes Utilizing the Bitcoin Blockchain." Future Generation Computer Systems, 2018.
URL https://doi.org/10.1016/j.future.2017.08.024

The documentation and verification of real-world events plays an important role in smart systems, e.g., with regard to supply chains or logistics. Once events have been identified, it is necessary to distribute them to data stakeholders using a trusted channel. Especially in distributed scenarios, where different organizations might be involved, it is necessary to store this data in an immutable way.

In the world of cryptocurrencies like Bitcoin, a similar problem arises, since transactions need to be stored in a permanent and unchangeable way without relying on a trusted third party. For this, the Blockchain is applied. Within this talk, we will outline how Blockchains can be used in order to provide runtime documentation and verification in business processes. In particular, we will present how the Bitcoin Blockchain is used in a concrete solution to achieve runtime verification for distributed, choreography-based business process execution.

4 Working Groups

In this section, we summarize the main results achieved by selected working groups.

4.1 Bridging the Gap between IoT and Blockchain: Research Questions & Challenges

Fabiana Fournier (IBM – Haifa, IL), Agnes Koschmider (KIT – Karlsruher Institut für Technologie, DE), Raimundas Matulevičius (University of Tartu, EE), Sooyong Park (Sogang University – Seoul, KR), and Stefan Schulte (TU Wien, AT)

In short, the Internet of Things (IoT) is defined as the pervasion of business and private spaces with "a variety of things or objects ... [which] interact with each other and cooperate with their neighbors to reach common goals" [3], forming "an interconnected world-wide network based on sensory, communication, networking, and information processing technologies" [35].

Blockchain technologies have been named as a facilitator for the IoT with regard to a number of different issues [15, 53], e.g., the verification of IoT data and actions.

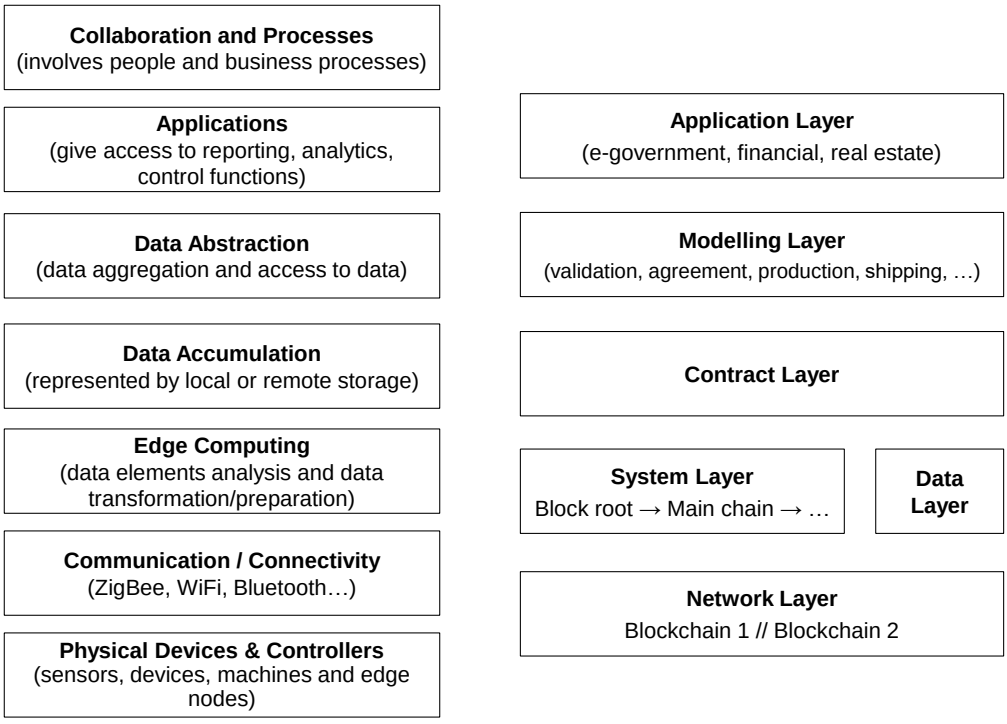

Figure 1 Layer Models for IoT (Left) and Blockchain (Right) [16, 70].

To categorize the application potential of blockchain technologies in the IoT and to identify research challenges, we make use of the IoT reference model introduced by Cisco [16], and the blockchain layer model introduced by Zhang and Jacobsen [70]. Figure 1 gives an overview of these two layer-based models.

In the following, we will discuss a number of particular research challenges with regard to the IoT and blockchain technologies. These comprise challenges where blockchain technologies may help to overcome IoT-related issues, as well as research questions which stem from the fact that IoT technologies and standard blockchains have partially contradicting features. For each research challenge, we describe where it is located in the combination of the two layer models.

4.1.1 Research Challenges

Storage capacity and latency: IoT nodes are generating data in high volume, high velocity, high variety (also known as the original "3 Vs" in Big Data [14]), and low veracity [29]. This is contradictory to some essential characteristics of blockchains: First, smart systems consisting of many IoT nodes may generate huge amounts of data, while data storage in blockchains is usually expensive, leading to the need to find ways to pre-filter data and/or store (parts of the) data off-chain in order to decrease the amount of data which needs to be stored on-chain. Second, the velocity of IoT data makes it necessary to process it fast (especially in data stream processing [27]), while the generation of blocks in blockchains leads to an inherent delay. Therefore, blockchain technologies may be used for other purposes, but not necessarily for real-time data exchange in the IoT. Third,

data variety makes it necessary to provide solutions that different data types can be used in blockchain-based systems. Finally, IoT data may lead to issues regarding data veracity, making it necessary to change stored data. However, blockchains rely on the principle that data cannot be changed. Hence, new solutions, e.g., for data versioning in the blockchain, need to be provided.

With regard to Figure 1, these research challenges can be attributed to both the *Data Abstraction* and *Communication / Connectivity* layers in the IoT reference model, and the *Data Layer* and *Network Layer* in the blockchain layer model.

Scalability: IoT-based smart systems can scale up and down very quickly, i.e., the number of nodes (and therefore the amount of data) in the IoT may vary a lot over time. However, some blockchain technologies, e.g., the Bitcoin blockchain, have trouble with scalability, since they are restricted to a certain amount of transactions per time frame. Therefore, it is important to select a sufficient blockchain technology for IoT scenarios. In some cases, the usage of side-chains might be a good advise.

With regard to Figure 1, this research challenge can be attributed to the *System Layer* in the blockchain layer model.

Security: Security is an issue on all levels of the IoT stack. Blockchains may be able to help to take care of some of the issues here, e.g., to identify nodes and entities in general, based on the keys already used in blockchains. Permissioned blockchains may be used to establish closed IoT systems, where data is only shared and forwarded between known partners and nodes, thus decreasing security issues.

With regard to Figure 1, this is an overarching research challenge, which needs to be regarded on all layers in both the IoT reference model and the blockchain layer model.

Anonymity and data privacy: Per se, data which is stored on a blockchain is available to all parties which can access the blockchain. Even though parties are anonymized (or rather pseudonymized) through their blockchain addresses (as in Bitcoin), their identities might be derived from the stored data. Therefore, data models are necessary which can be used in order to categorize which kind of data can be stored in which kind of blockchain, or which data should only be stored offline.

With regard to Figure 1, this is an overarching research challenge, which needs to be regarded on all layers in both the IoT reference model and the blockchain layer model. However, naturally, the *Data Layer* in the blockchain layer model and the *Data Abstraction* and *Data Accumulation* layers in the IoT reference model are of primary interest.

Smart contracts: With regard to IoT technologies, there are some examples for simple Event-Condition-Action smart contracts, e.g., to release funds once an (IoT-equipped) shipment has been delivered, e.g., [34], but there may be potential for further types of smart contracts in the IoT. In addition, data velocity may play a role here: If an oracle in a smart contract is based on IoT data, which is outdated very fast, this may lead to inconsistencies and conflicts, so that the execution of a smart contract cannot be validated. Some applications of IoT require timing constraints (e.g., fire alarm starts within 0.1 sec if temperature over 80°C in a smart farm). If this situation is controlled by a smart contract, there is a need for a real-time smart contract that satisfies timing constraints.

With regard to Figure 1, this is a research challenge closely related to the *Contract Layer* in the blockchain layer model.

Legal issues: While not IoT-specific, the legal status of smart contracts and blockchain-hosted data is not clear in all countries. In general, IoT data underlies data privacy legislation, as has been discussed above. Hence, data privacy legislation like the EU's

General Data Protection Regulation needs to be taken into account when storing data in an append-only store like the blockchain.

While not a legal issue, blockchains may also help to overcome legal disputes, by providing data documenting or verifying a particular situation (e.g., in autonomous driving [19]).

With regard to Figure 1, this research challenge can be attributed to the *Applications* layer in the IoT reference model, and the *Application Layer* in the blockchain layer model.

Consensus: While there might be powerful edge devices in the IoT, a very large percentage of all IoT devices only possess minor amounts of computational resources. Hence, mining and taking part in consensus finding on these devices is very difficult to achieve, especially in Proof of Work (PoW)-based blockchains. This makes it necessary to come up with solutions where IoT devices use a proxy to participate in a blockchain network, or are decoupled from a blockchain in another way. Also, blockchains which do not make use of PoW could be beneficial in IoT scenarios, since it might not be necessary to find consensus (but to validate new blocks) in the first place. Also, more lightweight consensus algorithms could be applied in the first place.

With regard to Figure 1, this research challenge can be attributed to the *System Layer* in the blockchain layer model and the *Physical Devices & Controllers* layer in the IoT reference model.

Low computing power and network issues: As already pointed out above, IoT devices may not be powerful enough to participate in certain blockchain-related tasks. This calls for lightweight consensus algorithms, abstract block mechanisms, lightweight hash functions, and also a minimum communication protocol.

IoT devices might be battery-powered, which makes it necessary to run them in an energy-efficient way. Thus, such devices cannot be "always-on". Also, data transfer between a blockchain-based backbone and the IoT devices needs to be limited – not only because of the energy consumption in the case of data transmissions, but also, since the network might actually become a bottleneck in large-scale IoT-based systems, where a huge number of entities are sending and receiving data. Together with the already mentioned consensus-related issues, this calls for blockchain solutions where IoT devices only join the network in particular cases, and are more or less dormant at all other times. This avoids that the blockchain-inherent communication overhead leads to rapid energy consumption.

With regard to Figure 1, this research challenge can be attributed to the *System Layer* in the blockchain layer model and the *Physical Devices & Controllers* and *Communication/ Connectivity* layers in the IoT reference model.

4.2 Two Perspectives on Blockchains: Capabilities vs. Features

Søren Debois (IT University of Copenhagen, DK), Marlon Dumas (University of Tartu, EE), Stephan Haarmann (Hasso Plattner Institut, DE), Hans-Arno Jacobsen (TU München, DE), Mieke Jans (Hasselt University, BE), Jan Mendling (Wirtschaftsuniversität Wien, AT), Mark Staples (Data61, CSIRO –Sydney, AU), Barbara Weber (Technical University of Denmark – Lyngby, DK), Francesca Zerbato (University of Verona, IT), and Kaiwen Zhang (ETS – Montreal, CA)

Blockchain technology is the subject of substantial enthusiasm and notable financial successes. For example in June 2018 alone, almost USD$6B worth of tokens were issued in ICOs [1].

Indications of widespread use of blockchain and distributed ledger technologies outside of tokens and cryptocurrency are emerging.

Prior work proposed different decision models with the goal to help answering the question "Do i need a blockchain for my application?" [21, 26, 39, 41, 42, 44, 50, 67, 68, 71]. Moreover, there are proposals to guide the design process (i.e., which blockchain configuration to best choose) (e.g., [68]). However, there is little work up to now that focuses on the business capabilities that might form part of a blockchain-based application supporting business operations and on how they link to blockchain features.

In the remainder, we proceed as follows. In Sect. 4.2.1 we provide the foundations of blockchain and distributed ledger technologies. Then, in Sect. 4.2.3 we give a list of business capabilities identified as key to blockchain. Moreover, in Sect. 4.2.4 we identify a list of blockchain features. Subsequently, we map features to blockchain capabilities. Finally, we outline potential future work in Sect. 4.2.5.

4.2.1 Background

In this section, we recall the main concepts of blockchain and distributed ledger technologies.

Blockchain and Distributed Ledger Technologies

Applications of blockchain typically shift trust from a third party (a bank, a government institution, a credit card company) onto something else, typically the technology of the chain itself. There are two reasons one might desire such a shift:

1. One does not wish to trust the third party. (Bitcoin: government-less money)
2. The third party is expensive. (Hypothetical example: Credit card companies.)

However, new applications might arise where there previously were no solution because involved parties could not or would not agree on a trusted third party. For example, Mærsk and other shipping companies always had the option of developing a global, centralised repository of shipping documents; however, presumably, who would control that repository prevented it from coming into existence.

We emphasise that in the absence of risk or trust issues, a blockchain has no purpose. In other words, *a blockchain is needed only if the data consumers and the data owner are in separate trust domains and the consumer has high-integrity requirements. There is no need*

[1] https://www.coinschedule.com/stats.html

for a blockchain when the data consumer(s) and data owner are in the same trust domain (e.g. inside a company).

To understand what capabilities are central to / indicative of such shifting of trust, we first (attempt) to clarify what is "trust" and what is "a capability".

Definition of trust. Trust is the acceptance of risk. Such risk may arise either from, say, malicious intent, or unintentional byzantine errors (either because of incompetence or because of hostile environment)

Alternative viewpoint (solution-driven).

The benefits conferred from blockchain technology constitutes "affordances" (see: Gibson, J.J., The Ecological Approach to Visual Perception. 1979. Boston: Houghton-Mifflin) rather than a outright features:

- Having trust in a system without having a trusted third party;
- Lower cost for the service;
- Lower barrier of entry;
- More accessible than traditional services (strategic advantage);
- Elimination of TTP;
- Tolerance to failures (impact of failures).

4.2.2 Capabilities and the Resource-Based View of the Firm

Business each have a wide range of capabilities [2]. Some are strategic capabilities which are key to the business's sustainable competitive advantage, and are valuable and distinctive compared to other businesses. Strategic capabilities are sometimes called "core competencies". Others are operational capabilities, which are necessary for the operation of the business, but will not be distinctive, and are more likely to be outsourced. A capability area may be strategic for one business, but operational for another.

Blockchains provide a mechanism allowing businesses to shift trust within the operation of their ecosystems. Often this is for disintermediation, stopping the centralised control of that capability by those third parties. This can be good for businesses that want to use that capability as an operational capability. However for trusted third-parties, this capability is a strategic capability, and blockchain may directly undermine the sustainability of their competitive advantage from that capability.

Definition of capability. "Capability thinking also means being aware of in what context the enterprise has the capacity and ability to offer business services that contribute to achieving business goals. The context basically captures what legal, technical, process, content, or other situation the business service is prepared for and what variations in providing the business service apply for what situation" [54].

4.2.3 Capabilities of Blockchain-Based Systems

Table 1 shows the main business capabilities resulting from our analysis and discussion. The list of business capabilities we have identified below is not exhaustive, and the capabilities

[2] Here we do not mean "object capabilities" which are secure references used in capability-based security models. We also do not mean software engineering capabilities captured for example in models such as CMMI.

Table 1 Business capabilities for blockchain-based systems.

BC1: Voting - Anonymous voting - Delegatable voting (conditional voting with smart contracts) - Number of participants(N): un/bounded - Non-sellable	Entry of votes submitted by different parties, tallying, and announcement of results.
BC2: Payment - Anonymous payments - Escrow payments - Variable payments - Complex conditional payment	Transfer of cryptocurrency between different parties.
BC3: Asset transfer	The transfer of assets (cryptocurrency, tokens) from one party to another.
BC4: Settlement (payment vs. delivery)	Synchronisation of simultaneous asset transfers.
BC5: Exchanges	Settlement of particular assets.
BC6: Introductions	Connecting parties interested in being endpoints of contacts
BC7: Referrals	Introductions where one or more party must be endorsed, authorised, and or made aware of by another.
BC8: Reputation	The reputation is a global score for participants representing trustworthiness.
BC9: Bookkeeping	Recording of transactions, typically for the purposes of financial reporting.
BC10: Brokering	Introductions for asset-transfer contracts.
BC11: Monitoring	The automated detection of transactions or contract executions satisfying particular, predefined properties.
BC12: Offering (incl. auctions)	Contract / transaction with initially undetermined counterparty.

may be interrelated (for example, settlement will involve payment). In addition, the business capabilities we have focussed on are multi-party capabilities, rather than capabilities that are mainly internal to a company.

4.2.4 Features of Blockchain-Based Systems

We then identified a list of blockchain features (system capabilities) as outlined in Table 2.

We then mapped the different business capabilities to the corresponding blockchain features (cf. Table 3). Optional features are listed in brackets.

A combination of the above described capabilities can be used to form a market.

Additionally, we identified the following inter-dependencies among features.

1. Audit Trail → Transactions → Signature → Encryption → Wallet information;
2. Contract → states → Verification;

Table 2 Features of blockchain-based systems.

F1: Data access on-chain	Storage; universal access to data stored on the ledger for any processing node.
F2: Encryption	Ability to encrypt and decrypt data stored on the blockchain.
F3a: Channel	Need-to-know access to data. Access control list.
F3b. Vault/Wallet information	Access to private information necessary to operate on the blockchain, but should remain confidential (e.g. private keys).
F4: States	Ability to record state for assets defined on the ledger, and transition the states using smart contract executions.
F5: Audit trail	Ability to record and link events in a sequence (provenance, logging, states are chained).
F5b: Receipts	Ability to obtain a detailed record per transaction, indicating which assets were read and modified.
F6: Transactions	Ability to submit transactions.
F6b. Permissions to submit data on-chain. F7. Identity management	
F8: Contract	Ability to invoke programs through transactions, and store contracts on-chain.
F9: Process	
F10. Verification	Integrity check of the ledger, and contract execution.
F11: Time service	Authoritative source of physical time, and timestamping.
F12: Notary service	Ability to put trust in / responsibility for a particular computation step in a given participant.
F13: Oracles	Special case of notary which injects external information into the system. (A mechanism for ensuring integrity of data provided transparently by a trusted data source.)
F14: Tokens	
F15: Anonymization	
F15b: Pseudonymization	
F16: Watermarking	Ability to permanently fix a signature inside a document stored on the chain.
F17: Digital signature	Ability to attach a signature to a transaction / document on the chain.
F18: Event	Ability to send events between accounts, to trigger smart contract invocations, and to notify external subscribers.
F19: What-if analysis	Ability to query the projected impact of a transaction / contract execution on the current state of the blockchain [9].

Table 3 Mapping Capabilities to Features.

Capabilities	Features
Voting	Transaction, Time service, (Anonymization, Notary, Identity Management, Tokens)
Payment	Transactions, Receipts, (Channel, Time service, Tokens)
Asset transfer	Transactions, Tokens, Watermarking, (Channel)
Settlement	Audit trail, Tokens, Notary, Contract
Exchanges	Transactions, Tokens, Assets transfer, Notary
Introductions	Process, Data access, Channel
Referrals	Transactions, Tokens, Identity Management
Reputation	Identity Management, Audit Trails, (Oracles)
Bookkeeping	Audit trails, Receipts, States
Brokering	Identity, Contract, Transactions, State, What-if
Monitoring	Audit trail, Events, Process, Contract (Time Service)
Offering (incl. auctions)	Transaction, Contract, Digital signature, (Time service, pseudonymization)

3. Time service → oracle → Notary;
4. Channel → Identity management → Encryption;
5. Tokens → Transactions.

4.2.5 Conclusion

This summary has taken initial steps towards identifying both the features that can reasonably expect to be supplied by a blockchain platform on the one hand; and the capabilities which applications for that platform may require on the other.

In the future we would like to investigate which features are supported by different blockchain platforms, to guide the decision which platform to choose.

Moreover, as another avenue of research we might look into different solution patterns on how to implement different features.

4.3 Factors Influencing Process Analytics on Distributed Ledgers

Claudio Di Ciccio (Wirtschaftsuniversität Wien, AT), Luciano García-Bañuelos (University of Tartu, EE), Mieke Jans (Hasselt University, BE), Jan Mendling (Wirtschaftsuniversität Wien, AT), Petr Novotny (IBM TJ Watson Research Center – Yorktown Heights, US), Ludwig Stage (Tübingen, DE)

Blockchains trace the sequence of tasks carried out in the course of business process executions by the totally ordered recording of transactions between involved parties, and additionally the logs of events registered by Smart Contracts. This leaves ample room for the ex-post analysis of conducted operations, for analytics, auditing, and mining purposes [40]. However,

it also poses questions on how to design the data storage, keeping into account a.o. the following facts: firstly, the recording of information, or the absence thereof from the stored data, influences the knowledge that can be extracted from ledgers; secondly, the exchange of data in blockchains such as Ethereum is expensive both for what the transactions fees and the write operations from Smart Contracts are concerned [66]; thirdly, saving information as data values carried as transactions payload entails a better traceability, on the one hand, but also unlimited access to the possibly sensible information exchanged, which is detrimental from a privacy viewpoint, on the other hand; finally, multiple blockchains can be adopted that can be either homogeneous or heterogeneous, e.g., a federated network of Ethereum ledgers (e.g. Quorum[1]) for the general inter-organisational, multi-party collaboration, and a federated Hyperledger Fabric for sub-processes involving sub-groups among the participants. On top of this, the introduction of querying languages for data stored as transactional or logging information plays a pivotal role in the introduction of process analytics over blockchains. To that extent, the preliminary contributions provided by the state query languages of Hyperledger Iroha [2], Hyperledger Burrow [3], and the querying of the backing MongoDB database through EOS [4], provide promising results.

In this report, we focus on challenges and requirements for conducting business process analytics on data stored by blockchain-backed process management systems. In particular, we examine the cases in which information is stored fully on-chain or partially off-chain.

4.3.1 On-chain challenges and requirements

In this section we investigate the case in which all the information that is relevant to the process execution is stored on-chain. Discussions on data management and provenance associated with this strategy, including the privacy concerns and the transaction costs, go beyond the scope of this summary. Even under the assumption that the issues related to those topics were appropriately handled, we envisage in the following some crucial aspects to be taken into account for the analysis of this data.

Audit-completeness

Starting with the fully on-chain, single ledger design, the audit-completeness of this data is paramount. Taking inspiration from a requirement set by process mining, if criteria are not provided to uniquely identify the transactions pertaining to a process instance, then linking the evolution of the process becomes a hard manual work at best, thus hampering the process analytics endeavours [6]. In the context of financial auditing, the auditor needs to consider both relevance and reliability of audit evidence. In the context of using blockchain technology, both dimensions might be impacted. The data that is stored on the blockchain ideally contributes to financial reporting assurance (relevant data). Further, providing evidence that data is sufficiently reliable, is often challenging to the auditor [46]. When evaluating this aspect, accuracy and completeness of the data are considered; two aspects that may be impacted positively impacted by blockchain.

[1] https://www.jpmorgan.com/quorum
[2] https://www.hyperledger.org/projects/iroha
[3] https://www.hyperledger.org/category/hyperledger-burrow
[4] https://eos.io/

Eventual consistency

As explained in the CAP theorem [10], distributed systems can enjoy at most two properties out of Consistency (every read receives the most recent write or an error), Availability (every request receives a non-error response), and Partition tolerance (the system continues to operate despite an arbitrary number of messages being dropped or delayed by the network). Reportedly, for instance, Ethereum does not guarantee (strong models of) consistency [4], but only eventual consistency [64]. This signifies that the monitoring of transactions carried out on a local node does not guarantee full reliability. We identify in this context the following *audit patterns of deviation*:

Reordering Transactions, as well as the data they bring, could occur re-ordered in case the local world state in a node gets changed by the substitution of the latest block, or a sub-chain, with a fork that achieved a larger consensus;

Recurring Supposing that a fork lead to a local history rewriting, an already analysed block could be replaced with a new one in which a processed transaction does not occur any longer, yet it recurs in a new mined block thereafter; in such a case, the same information might be included twice if a consistency check is not operated that rearranges the parsed information accordingly;

Missing In the case of forks, transactions that were considered as valid could be excluded by the agreed-upon fork, and then not re-included in case new transactions mined in the new blocks make them invalid; this poses the challenge on whether to discard the retrieved information when the corresponding transaction gets erased from the blockchain.

Abstraction and reverse engineering

In Ethereum, information stored on-chain can occur as event logs emitted from Smart Contracts or as payload to transactions, aside of the internal state of contracts which is however not explicitly written on transaction receipts but has to be recomputed by executing the code locally to the nodes in order to be undisclosed. Event logs and data parameters of the transaction can reveal explicit notifications and context specifications respectively, upon deserialisation[5]. Nevertheless, the way in which logs and exchanged data are engineered is tightly bound to how the the Smart Contracts are encoded. This hampers the ex-post interpretation of those sources of information, let alone their automated analysis. The promised verification and traceability of executed processes ends up being ad-hoc, and demanding manual effort, not so differently from what used to happen when striving to understand the behaviour of legacy systems through their logs [47]. This calls for the introduction of a specification language, possibly using the decorator pattern [22], to enrich the code of methods in Smart Contracts for the sake of self-documentation about state, data, and logging.

4.3.2 Off-chain challenges and requirements

In the remainder, we examine the case in which data are held out of the ledger. Again aside of the considerations on the data management and administration, we portray the envisioned challenges and requirements that involve the merging of operational information retrieved from the blockchain and the affected data from the outside.

[5] https://solidity.readthedocs.io/en/develop/abi-spec.html

Between ledger and the outer data

As the information is split between on-chain and off-chain data, the analysis necessitates of a mechanism to join at least two sources of information, one logging the sequence of actions mediated by methods invoked on smart contracts, the other reporting on the updates on, or retrieval of, data witnessing the conducted tasks. Technologies such as the InterPlanetary File Sysmte (IPFS) provide a mechanism for uniquely linking data chunks spread among peers outside the main blockchain. However, here we refer to those cases in which parts of the data pertaining to the process activities are kept in other systems that can be disconnected from the ledgers, such as external DBMSs, either centralised or federated. Especially in such a case, the association of ontologies to the process specification seems paramount to describe the semantic connections.

Versions of blockchain artifacts

Should the process undergo a redesign phase, the Smart Contracts implementing the old version could be replaced by newer ones. However, whilst the versioning of processes is a feature that is implemented by current Business Process Manaement Systems (BPMSs), the concept of contract replacement is not natively supported by blockchains such as Ethereum. Fully on-chain software architectures such as the one of the blockchain-based BPMS Caterpillar [37] may cater for it, thanks to their implementation of the factory pattern for generating Smart Contracts. However, in partial on-/off-chain scenarios, keeping track of the changes entailed by subsequent versions of the involved artifacts becomes a challenge of higher difficulty, as it involves at once information integration and object matching, together with the semantic version control over software and data updates.

Between ledger and reality

Solutions to connect the digital world of the blockchain with outer reality are crucial to cater for business processes interacting with physical objects, such as in the case of the manufacturing, logistics, or healthcare domains, to mention but a few. For instance, the notion of time is implemented on blockchains such as Ethereum as *block-time*. Such a timestamp is shared among all transactions therein, thus at a coarse-grain level. The aforementioned business processes operate in real time instead. To solve this problem and inject information on real-world information including time, so-called *in-bound oracles* such as Oraclize[6] have been introduced. Oracles operate as a middleware connecting reality with the on-chain information space, and take on the task to return consistent answers to virtually all nodes in the bloakchain that execute the same Smart Contract locally. The out-bound connection seems however more challenging. Owing to the concept of eventual consensus, an operation executed by the run of a Smart Contract may be withdrawn, should a different suffix of the blockchain eventually reach consensus over the local version. From a BPM perspective though, the execution of certain tasks should not be subject to rollback, especially if it leads to permanent changes in the real world. Waiting times could be introduced on purpose between the local transaction and the actual execution of the associated operation in real world, so as to reduce the probability that the task is undone. This is indeed what happens with many purchases paid in Bitcoin. However, this approach might lead to delays that encumber or disrupt the executions of business process in general. Besides, only mitigating the uncertainty of task executions is not sufficient in many cases. An approach that eliminates the risk of operation rollback is thus of high relevance and calls for future research endeavours.

[6] https://github.com/oraclize

4.4 A Holistic Vision of Blockchain-Based Application Design, Specification, and Implementation

Michael Coblenz (Carnegie Mellon University – Pittsburgh, US), Richard Hull (IBM TJ Watson Research Center – Yorktown Heights, US), Qinghua Lu (Data61, CSIRO – Sydney, AU), Ingo Weber (Data61, CSIRO – Sydney, AU)

Blockchain is an emerging technology, and we are just at the starting point in understanding its potential and where it can bring real value. The existing design and specification methods are suboptimal for blockchain-based applications, and not well integrated – e.g., BPMN does not support data well enough for many use cases on blockchain; on-chain and off-chain parts of the application are viewed disparately (and this separation is unlikely to resolve, due to the nature of some data and computation being required to remain off-chain); and governance for the evolution of blockchain smart contracts brings challenges not found in other contexts. Our vision is to have a holistic design/specification/implementation approach which covers various disparate components in an integrated fashion.

4.4.1 Problem Statement

Background: Blockchain is a new distributed ledger technology, which has attracted broad interest, including in industry and government, in exploring how to use blockchain to re-design systems and to build the next generation of applications. Issues include the following:

1. The existing design and specification approaches are suboptimal and not well integrated. First, business analysts usually lose control of system development once they complete requirement specification since they can hardly understand the smart contract programming language and check the conformance with the requirements they specify. It is also hard for domain experts to inspect the code to understand how their ideas are represented in the system. Second, designers need to work on different types of model abstractions which are not well integrated. Third, blockchain is by design data-centric, which can hardly be supported by existing models, e.g. BPMN does not support data modeling well.
2. The architecture of blockchain-based system is usually separated into on-chain and off-chain parts since blockchain has limited storage capability and the information on blockchain is designed to be accessible to all the participants. This separation of on-chain and off-chain is unlikely to be resolved completely, due to the nature of some sensitive and large sized data, compliance requirements, and some computation being required to remain off-chain (in part due to the high cost of on-chain computation).
3. Blockchain systems raise challenges in governance not typical in other domains. These systems will increasingly be designed and maintained through collaborations of multiple stakeholders coming from different organizations. Agreed upon smart contracts are themselves stored on the Blockchain to support increased trust in the system. Furthermore, because the data is immutable, there is an increased desire for reverse compatibility of smart contract versions. This to enable easier usage of data from earlier versions of a smart contract, and in auditing and analytics solutions layered on top of the Blockchain.

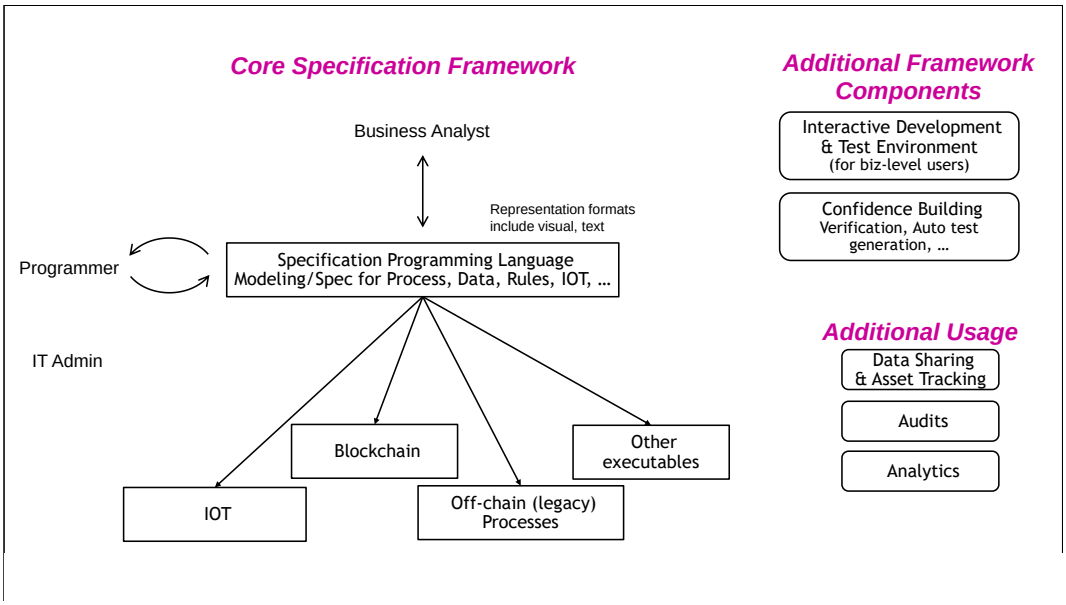

Figure 2 High-level framework for design, specification, and implementation of Blockchain-enabled solutions.

4.4.2 Proposed Approach

Our vision is to have a holistic approach for the design, specification, and implementation of collaborative information systems using blockchain, which covers various disparate components in an integrated fashion. For instance, there might be a set of models for static and dynamic aspects (e.g., data, process, business rules, required inputs from user and devices) which is translated into executables for blockchain (smart contracts), enterprise systems (including DBMS schemata and internal processes / rules), IoT device instructions, and UIs. By starting from the beginning, we can leverage the characteristics of the blockchain. We propose a technique that uses research methods in human-centered design in order to maximize the chances that the resulting system will be as effective as possible for its users. By doing so, we will support different roles, including business analysts, domain experts, and programmers, with one holistic method. Reasons for this requirement include the need of business users to own and understand the application well enough so that they can ensure that the resulting implementation conforms to their expectations, and they can assess risks, compliance, and likely or probable effects.

4.4.3 Research Questions

User-oriented questions

A solution should address the following general categories of users, representing several different collaborating organizations:

- Business analysts
- Programmer/software engineers
- IT administrators
- End users
- Lawyers

In order to assess the needs of the users, we should address the following questions:
- What use cases do systems need to support for each user?
- How do users expect to express solutions to their problems in their own domains?
- In order to make tangible progress in the shorter term, what modeling paradigms (e.g. BPMN, CMMN, etc.) can be adapted and integrated to form a first workable programming model/language?
- What principles can guide creation of usable integrated development environments for business analysts?
- How can we characterize the domain of applications that need to be supported? Is existing work in this area sufficient? Academic work on blockchain may not capture enough of the business use cases, and industrial efforts to date may not be publicly accessible or represent the whole range of possible use cases for the tool set.

Multiple target platforms

An individual blockchain application may comprise components for more than one platform, such as:
- Blockchain programs
- IoT
- Off-blockchain server-side programs
- User-facing systems (e.g. the UI for a dapp)

This leads to the following research questions:
- How can a system facilitate interoperation among the different components? For example, data structures may be passed among different components; we want to ensure consistent semantics for a given object/asset.
- To what extent can portions of the application be targeted at executables in a flexible way? For example, an initial prototype might run entirely off-chain, with components moving to the blockchain as needed or as scalability is demonstrated. Or an application might assume that many IoT devices are going to provide data, and later the data may be sourced via traditional transaction invocations.
- Do the different targets need distinct kinds of specifications/implementations? Is it appropriate to consider UI as part of the blockchain application, or is it more properly considered as a separate project that invokes APIs?

Underlying blockchain platforms

What are the relevant properties of blockchain platforms that vary among platforms (e.g. Hyperledger Fabric, Ethereum, etc.)? How are applications customized/optimized to run on particular platforms? For example, how should an application designer/implementer decide how to represent the state need to be serialized to the ledger?

How do governance needs affect application development needs? For instance, how can blockchain applications evolve over time?

Modeling-oriented questions

The model-driven community tends to assume that it is best to separate a model of the system from the implementation, which may be partly or completely generated from the model. However, this separation might unnecessarily add complexity to the system (requiring some or all users to understand two different perspectives of the system and keep them in

sync). Further, it might result in mismatch between people's expectations (in the model) and the reality of the lower-level implementation.

- To what extent should the users work on the implementation directly versus work on several different abstractions of the system?
- What kinds of views facilitate the reasoning that different stakeholders need? Which of these views need to be editable, and in what cases might higher-level changes accidentally invalidate previous lower-level modifications?
- To what extent can we re-use existing modeling paradigms (e.g., BPMN, CMNN, UML class diagrams, SBVR, etc.) vs. develop variants and or new paradigms? If some existing modeling paradigms are used, how can they be combined together (e.g., how to combine the process-centric perspective of BPMN with the data-centric perspective of UML class diagrams?)

Analysis/auditing

Blockchain provides the opportunity to hold a shared global view of all data relevant to a collaboration between organizations. This data can be used in ways that enterprise-specific data has been used in the past, including to track assets, support audits, and perform analytics. This raises several research questions, including the following.

- What use cases should auditing tools support? What are the needs of the users, and who are they?
- What design principles should be followed with regards to data and processing, in order to enable conceptually simple asset tracking, auditing, and analytics? This should permit efficient implementation on various blockchain technologies, with regards to execution of smart contracts and of the data processing for tracking, auditing and analytics.
- Some blockchain technologies include privacy and access controls concerning both data and processing steps. How do these interact with support for tracking, auditing, and analytics?

Evaluation

We did not design a particular evaluation plan. An evaluation plan would need to address these questions:

- What kind of user studies would provide evidence regarding usability from the various perspectives?
- What case studies should be done to evaluate expressiveness and to help iterate on the system design?
- Are there formal properties of the system that should be verified?

4.4.4 Summary

This working group discussed a vision of a holistic approach for design, specification, and implementation of blockchain-based information systems. The main part of the discussion focused on research questions, covering user-oriented aspects, target platforms across blockchain and off-chain components, IoT, and UIs, blockchain platforms, modeling, analysis, and evaluation aspects. Achieving this vision would require major effort, but could yield many benefits over existing solutions.

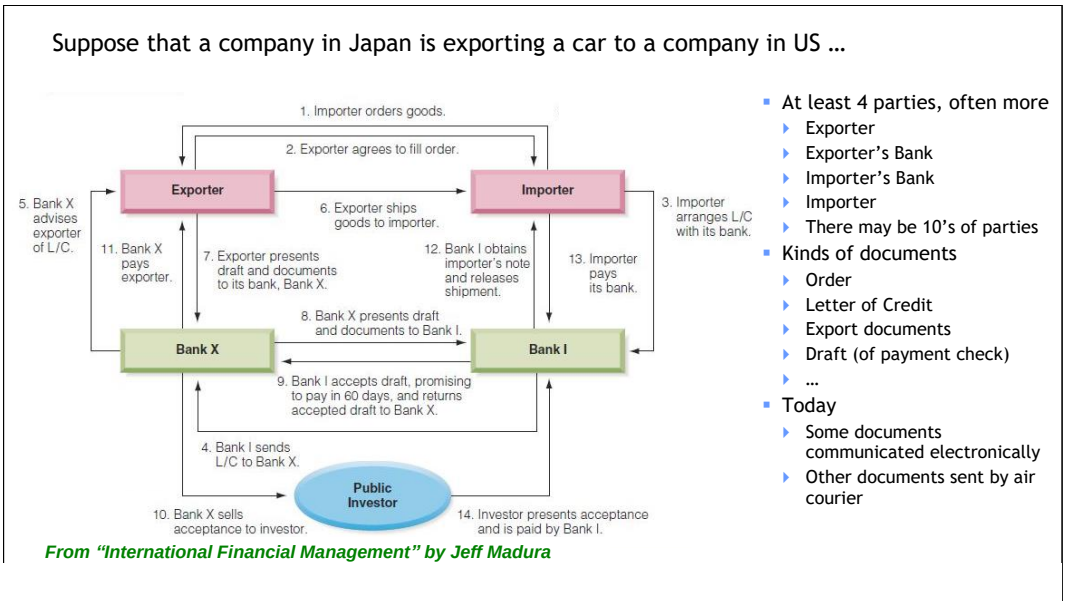

Figure 3 Simplified view of Trade Logistics use case.

4.5 Towards a Blockchain Collaboration Meta-Model and Language

Michael Coblenz (Carnegie Mellon University – Pittsburgh, US), Claudio Di Ciccio (Wirtschaftsuniversität Wien, AT), Marlon Dumas (University of Tartu, EE), Fabiana Fournier, Luciano García-Bañuelos (University of Tartu, EE), Richard Hull (IBM TJ Watson Research Center – Yorktown Heights, US), Qinghua Lu (Data61, CSIRO – Sydney, AU), Raimundas Matulevičius (University of Tartu, EE), Jérôme Siméon (Clause Inc. – New York, US), Mark Staples (Data61, CSIRO – Eveleigh, AU), Ingo Weber (Data61, CSIRO – Sydney, AU)

A blockchain collaboration is a process involving multiple *participants* that interact between them by means of a blockchain. An execution of a collaboration (a *collaboration instance*) requires binding roles in the collaboration to particular participants. The participants create and manipulate a number of *concepts*. A *concept* is a container of data (possibly stored on the blockchain), such as Purchase Order, Invoice, Product, Letter of Credit, etc.

A concept is either an *asset* when it is subject to a (linear) ownership relation with the possibility of ownership transfer, or a *business object* otherwise. An *asset* has a number of properties, which may be attributes (e.g. the amount of an invoice), asset-to-asset relations, or asset-to-participant relations. An example of an asset is a house, a car, a parking place, or a product. A purchase order or an invoice are examples of business objects. Note that it is possible that a business object (not subject to ownership) can have a corresponding tradeable invoice, which is an asset and may thus be transferred. In this case, we distinguish the tradeable invoice as an asset and its corresponding invoice (business object) from which it originates.

Concepts may be fungible if they do not have an identity, or non-fungible in which case each instance of the concept has a unique identifier.

A business collaboration consists of a set of concepts, a set of participants, and a set

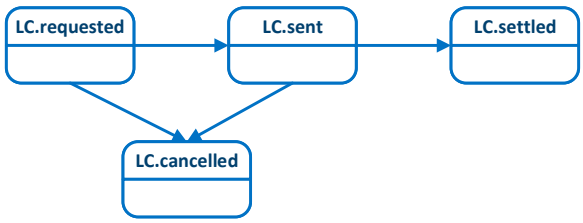

Figure 4 Letter of Credit lifecycle.

of transactions. A transaction modifies the state of the collaboration, which may imply modifying the properties/state of one or more concepts in the collaboration.

The occurrence of transactions may be constrained by the *collaboration's behavior*. At an abstract level, a collaboration behavior determines/restricts whether or not a given transaction can take place given the current state of the concepts involved in a collaboration.

The behavior of a collaboration may be defined at a local level (i.e. at the level of one asset at a time) or at a global level (i.e. at the level of a collaboration consisting of multiple assets). Below, we sketch multiple approaches for modeling collaboration behaviors both at a local and at a global level. To illustate each of these approaches, we make use of the trade logistics scenario depicted in Figure 3.

Three of the approaches are tied to a model-driven paradigm and include both visual and machine-readable specification (Sections 4.5.1, 4.5.2, 4.5.3). Two are focused on incorporating Blockchain-relevant constructs and abstractions into domain-specific programming languages (Sections 4.5.4), 4.5.5). There are other somewhat related approaches for orchestrating or choreographing the assets that are not considered here, e.g., the declarative style of object-centric behavioral constraints (OCBC) [61], the declarative choreographies for business artifacts [56], and the more procedural BPMN process choreographies [20]. Section 4.5.6 builds on some of these ideas to develop a framework for specifying collaborative logs in the context of Blockchain solutions.

4.5.1 A Statechart-Based Approach

In this approach, the behavior is modeled both at a local and at a global level using statecharts.

Each asset has a set of possible states, which may be related by means of a statechart diagram. For instance, a letter of credit has four states: "requested", "sent", "settled", and "cancelled" – with the lifecycle shown in Figure 4. Assets can also have different types of status (with a lifecycle each), and constraints can exist between them. For instance, a car can be in high-level states "active" and "inactive", and its loan status can be "owned outright" or "collateral to loan". In addition, a car may be in state "ready to ship". Transitions are labeled, e.g., a self-link from and to "owned outright" called "transfer ownership".

Because we are on a blockchain, there is no separation between Purchase Order and Sales Order; instead, the shared view between the parties is held on-chain, and has two associated roles: buyer and seller.

In addition to the local behavior (local statechart), we have one global state chart for the whole blockchain collaboration, which describes how the different lifecycles are connected to achieve the goal of the business collaboration. For instance, consider the global state chart in Figure 5, which proceeds as follows. The collaboration starts with an order being

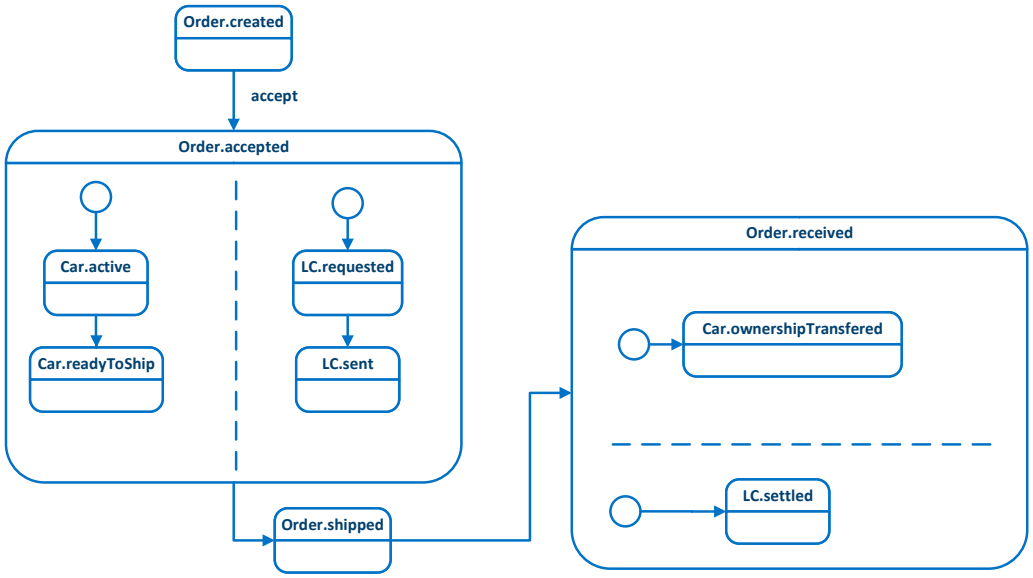

Figure 5 Global statechart.

created by the buyer, and subsequently accepted by the seller. Then, in parallel, the car is manufactured ("car.created") and in parallel the Letter of Credit (LC) is requested from the importer's bank ("LC.requested").

4.5.2 A BPMN-based Approach

In this approach, the behavior of the collaborative process is captured as a business process model in the standard BPMN notation. In this model, lanes represent Parties which collaborate together via a blockchain platform. The transactions that parties can execute on the blockchain are captured as activities in the BPMN process model. The behavior is thus captured via sequence flows, gateways, and other BPMN control-flow constructs.

The running example, captured using this approach, is shown in Figure 6 (specifically this model captures the collaboration between an exporter e.g. Toyota and an importer e.g. Auto Rip Off. The data objects manipulated in the process includes the Purchase Order (from Auto Rip Off's point of view) and the Sales Order (from Toyota's point of view). Transactions (Request order and Register sale) are written to ledger, which represented as a data store Order/Sales ledger in the model.

BPMN-based approaches have been proposed in the literature already, including [23, 65], and implemented in the tools Caterpillar [36] and Lorikeet [58]. Lorikeet has asset management and control features, which can be linked to process activities.

4.5.3 An approach inspired from Case Management

In this approach to the global orchestration across assets, parties, and their relationships, we focus on a modular model of *activities* along with a flexible, somewhat declarative approach for specifying when activities might be launched and concluded. In one sentence, the global structure of the activities follows the structure of a (hierarchical) Directed Acyclic Graph (DAG) where the control of activity launch and completion is controlled by a combination of events and conditions that refer to the global state. This follows the spirit of case modeling

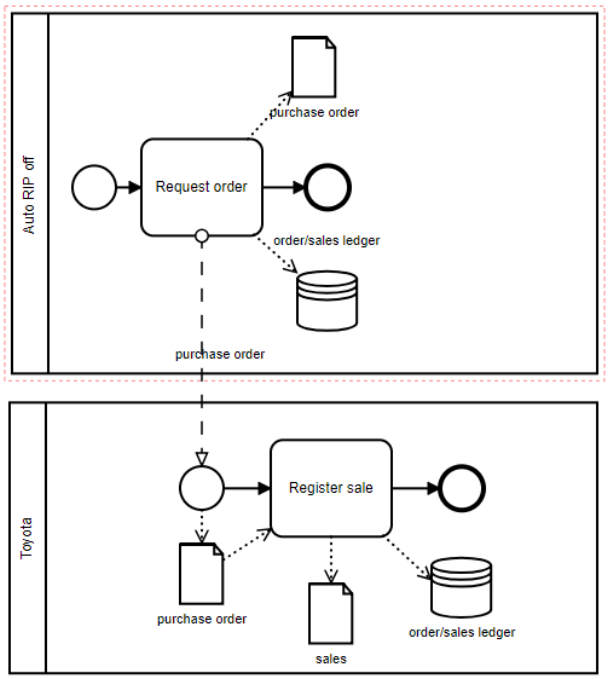

Figure 6 A BPMN-based process model.

in general (cf. OMG CMMN), the van der Aalst et al. Case Handling paper [62], and work on the Guard-Stage-Milestone (GSM) artifact-centric model [28, 18]. A formal operational semantics for this model can be developed along the lines of GSM, but would be much simpler because the focus is on DAG with rollback, rather than full GSM. A more detailed description and illustration now follows.

Figure 7 provides a high-level, simplified view of how the DAG with rollback approach might be visualized for the trade logistics example. The focus of this view is on the activities; the linkage to assets and their lifecycles, and to the relationships between assets and parties is not incorporated here. The notation follows that of CMMN, but with some variations.

The key building blocks are activities, including *atomic activity*, *long-running activity*, and *composite activity*. Activities are launched by events. The events might come from the outside world (i.e., transaction requests into the Blockchain), or may be *internal*, by which we mean triggered from the completion of some other activity in the orchestration. The activities may include *guards*, which are conditions on the state of the underlying assets and relationships, and also the state of the global orchestration. The solid arrows indicate events, including some internal ones, and the diamonds indicate guards. The diagram illustrates that the activities can be nested.

Rollbacks may occur, following the spirit of van der Aalst's Case Handling model, but extended to include modularity (and possibly some notion of compensating actions).

A key aspect of the DAG with rollback approach is the rich flexibility that can be incorporated with regards to the ways that activities might occur in sequence and in parallel.

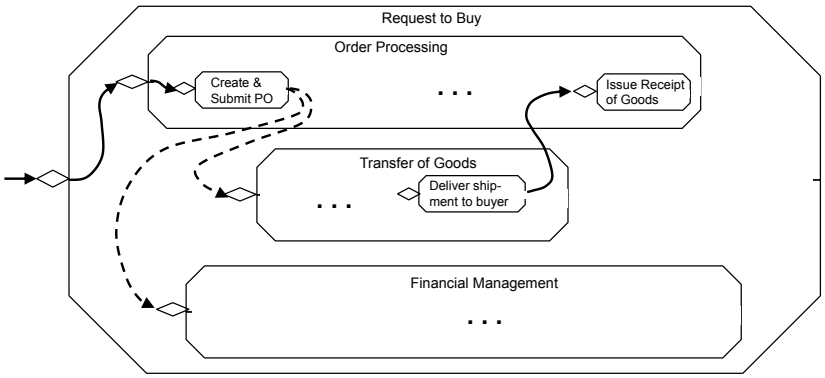

Figure 7 Illustration of DAG with rollback for global orchestration.

4.5.4 A legal contract perspective

Legal contracts are collaborative processes involving multiple participants, usually in a context where there is a measure but not entire trust between the parties. As such they are an interesting model, and an important use case, for Blockchain applications.

A legal contract typically comes with a natural language describing the terms of the agreement, its conditions and the process that is allowed or mandated between the parties. However, one can also consider a legal contract from an information system point of view. In this *legal contract perspective*, the contract is the central object being created and executed. This usually happens in three steps.

The first step is the negotiation phase during which the parties discuss the terms of the contracts, until they reach an agreement on those terms. From an information system point of view, this is similar to agreeing on the logic (or program) that will be executed.

The second step is to sign and *execute* the contract. From an information system point of view, this is similar to creating an instance of the logic (or program) that has been agreed upon.

The third step is contract execution, where the contract serves as an agent between the parties. During that step, the contract typically maintains a state and respond to requests (e.g., what is the price that should be paid for the goods delivered, including penalties if any).

Figure 8 gives an overview of this approach. In the top half of the figure is the contract instantiation, which involves creating a contract (e.g., from a contract template which has been negotiated and adapted for the specific parties and business needs). Once the contract has been instantiated, it can be executed. The execution involves sending requests to the contract, which returns a response and may change its state.

In this section, we present a purchase agreement between a car dealer and a car manufacturer based on the legal contract perspective. This example is loosely inspired by the one written in Obsidian that we saw in Section 4.5.5.

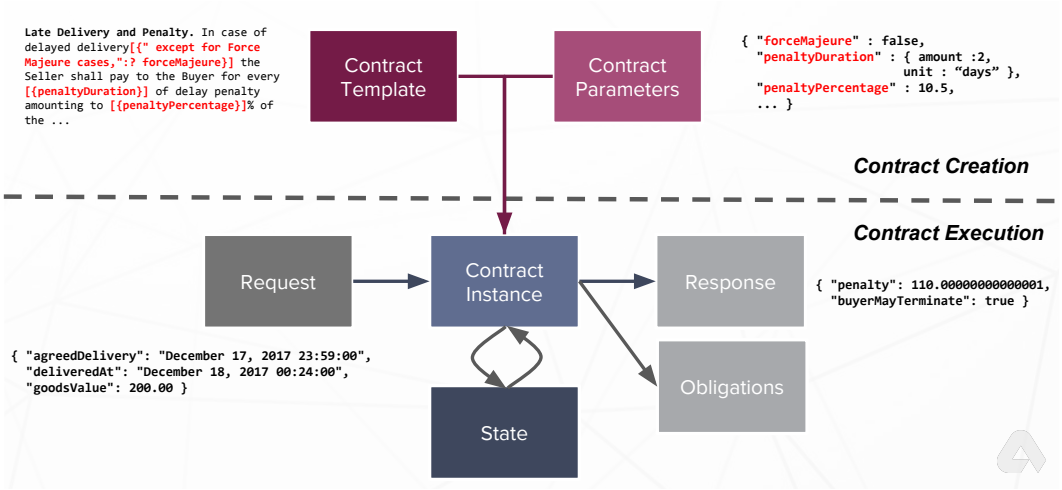

Figure 8 A Legal Contract Perspective.

Modeling the data

To model the data manipulated in a blockchain-based collaborative process, one can use the Composer Modeling Language (or CML) offered by the Hyperledger Composer platform.

This is a convenient way to describe a class hierarchy with some specific distinctions relevant to contracts (e.g., it distinguishes between general purpose `concepts`, `participants` which are the parties involved in the contract, and `assets` which can be owned and transfered between parties.

```
namespace org.dagstuhl.automotive

import org.accordproject.cicero.runtime.*

// Car and inventory
asset Car {
o String name
o Double purchasePrice
}

concept Inventory {
o Car car
}

// Participants
participant Dealer {
o String name
}

participant Shipper {
o String name
}

participant Manufacturer{
o String name
```

```
o Shipper shipper
o Inventory inventory
}

// Contract parameters
concept Parameters{
o Dealer buyer
o Manufacturer supplier
}

// Contract state
asset SupplyAgreementState extends AccordContractState {
o Car car
o Participant owner
}

// Contract transactions (request/response)
transaction Order extends Request{
o Car car
}
transaction Shipped extends Response{
o DateTime at
}

transaction Deliver extends Request{
}
transaction PaymentDue extends Response{
o Double amount
}

transaction Pay extends Request{
o Double amount
}
transaction Completed extends Response{
o String message
}
```

Note that the model includes parameters of the contract (i.e., which are the parties in this specific example), and the structure of the contract state (here which car is being purchased and which current participant owns the car).

This is only a simple example, one could model something much more complex, for instance the various parties here could be US Businesses with certain general information (where is it incorporated, etc). But for now this will do.

Describing the logic

The next step is to describe the logic of the contract. Here this is a very simple contract the contract is between a buyer (of type `Dealer`) and a supplier (of type `Manufacturer`).
There are only a few simple operations available in this contract:

1. The buyer can place an order and the contract responds with shipping information from the manufacturer
2. The shipper can deliver the car and the contract response with payment information

3. The buyer can pay the bill and the contract responds with a thank you note (or screams if the payment is not correct)

The following code for the contract is written in Ergo:

```
contract SupplyAgreement over Parameters state SupplyAgreementState{
// Intialize the contract
clause init(request:Request) {
set state SupplyAgreementState{
stateId: "IDLE",
car: contract.supplier.inventory.car,
owner: contract.supplier
};
return
}
// Place an order
clause order(request:Order) : Shipped {
enforce(request.car.name = contract.supplier.inventory.car.name);
set state SupplyAgreementState{
stateId: "INTRANSIT",
car: state.car,
owner: contract.supplier.shipper
};
return Shipped{ at: now() }
}
// Delivery clause
clause deliver(request:Deliver) : PaymentDue {
set state SupplyAgreementState{
stateId: "DELIVERED",
car: state.car,
owner: contract.supplier
};
return PaymentDue{ amount: state.car.purchasePrice }
}
// Payment clause
clause pay(request:Pay) : Completed {
enforce request.amount = state.car.purchasePrice;
set state SupplyAgreementState{
stateId: "COMPLETED",
car: state.car,
owner: contract.buyer
};
return Completed{
message: "Enjoy your " ++ state.car.name
}
}
}
```

As the reader can observe, there is a contract structure (akin to a class in an Object-Oriented programming sense), and a contract contains clauses (akin to a method in an Object-Oriented programming sense). There is a special init clause to set up the contract (akin to a constructor method in an Object-Oriented programming sense). In our example contract, we have one clause per operation available on the contract (one to order, one to deliver and one to pay).

In Ergo, the `return` statement indicates the response when calling a clause, the `set state` statement indicates a change of state in the contract, and the `enforce` statement is a precondition for the clause (i.e., the clause with return to the caller with an error if it is false).

Instantiating and invoking the clauses of the contract

Now that we have modeled our data and our contract, we are ready to instantiate the contract. This can be done with the following Ergo declarations:

```
// Let's create all the parties
define constant the_car = Car{
name : "Athena 360",
purchasePrice : 36000.00
}
define constant the_manufacturer = Manufacturer{
name : "AutomotiveInc",
shipper: Shipper{ name : "HappyShipping" },
inventory : Inventory{ car: the_car }
}
define constant the_dealer = Dealer{ name : "Best Deal Bros." }

// Now we can initialize one contract
set contract SupplyAgreement over Parameters{
buyer : the_dealer,
supplier : the_manufacturer
}
call init(Request{});
```

Now we are ready to ship cars! Here is an example of calling the contract from purchase order to completion. At each steps we show the response and the new state of the contract along with their types.

```
call order(Order{ car: the_car });
Response. Shipped{ at: dateTime("2018-08-16 16:16:42")} : Shipped
State. SupplyAgreementState{
stateId: "INTRANSIT",
car: Car{name: "Athena 360", purchasePrice: 36000.0},
owner: Shipper{name: "HappyShipping"}
} : SupplyAgreementState

call deliver(Deliver{});
Response. PaymentDue{amount: 36000.0} : PaymentDue
State. SupplyAgreementState{
stateId: "DELIVERED",
car: Car{name: "Athena 360", purchasePrice: 36000.0},
owner: Manufacturer{...}
} : SupplyAgreementState

call pay(Pay{ amount : 36000.00});
Response. Completed{message: "Enjoy your Athena 360!"} : Completed
State. SupplyAgreementState{
stateId: "COMPLETED",
car: Car{name: "Athena 360", purchasePrice: 36000.0},
```

```
owner: Dealer{name: "Best Deal Bros."}
} : SupplyAgreementState
```

4.5.5 Languages for direct editing

Some of the approaches in this document separate a meta-model from the programming language. In this section, we describe a different vision: choose one language but potentially provide visual editors. Such an editor might leverage some users' familiarity with prior meta-modeling tools, but with an important difference: the editor would edit the program directly rather than editing a model of the program, which might be semantically disconnected from the program. It might also facilitate code understanding and architectural analysis by providing diagrams that enable understanding of high-level behavior (rather than requiring readers to infer relationships between components).

By having a direct visual editor for the language, the program is always consistent with the visual representation, and there is the opportunity for gradual learning of the language. In meta-modeling systems, the user may be required to learn two different languages and edit them separately.

RequestToBuy: Roles, assets, and relationships example

This approach is intended to be directly representative of the approach we discussed, in terms of distinguishing roles, assets, and participants, mediated by transactions.

A role describes an entity that can have behavior, whereas assets are merely containers for data. Transactions can be invoked externally; in contrast, an `action` can be thought of as a step in a process, These declarations enable tools to infer high-level process structures (perhaps automatically drawing diagrams). See the top-level requestToBuy transaction, which declares that it invokes the processPurchase action on the manufacturer.

A `role` can be thought of as a kind of class or interface (this needs more thought) that captures relationships between participants. Likely, a participant will need to be able to fulfill more than one role. For example, ownership might be captured as a particular role.

This approach is sketched in the following listing.

```
role Dealing {

}

role Manufacturing { // gerund indicates role
Shipper shipper;
Inventory inventory;
}

participant Manufacturer fulfills Manufacturing {
// Placeholder: track fulfilled orders

action processOrder(Order o) {
if (inventory.contains(o.car)) {
Car@Owned car = inventory.remove(o.car);
s.ship(car);
// Placeholder: record order in list of fulfilled orders
}
```

```
}
}

participant AutoRipOff fulfills Dealing {
// ...
}

asset Car {
Manufacturing manufacturer;
// ...
}

asset Order {
int purchasePrice;
Car@Unowned car;
Dealing dealer;

Order(int p, Car@Unowned c, Dealing d) {
purchasePrice = p;
car = c;
dealer = d;
}
}

transaction requestToBuy (Car@Unowned car, int price)
returns Order // perhaps this should not  return an order.
by AutoRipOff d  // this implementation is specific to one dealer
issues processOrder to Manufacturer m
{
assert(car.manufacturer == m);
Order order = new Order(price, car, d)
manufacturer.processOrder(order);
return order;
}
```

RequestToBuy: Obsidian example

This is an Obsidian [17] implementation of the requestToBuy scenario. As of this writing, it compiles in Obsidian.

It is likely possible to automatically infer a high-level process diagram similar to Figure 5. One main discrepancy is that Obsidian has no notion of high-level "process"; such a notion might need to be added so that sequences of transactions could be given appropriate names for the diagram.

```
contract Shipper {
 ...
transaction ship(Car@Owned >> Unowned car) {
// Placeholder: take ownership of the car.
// For now, throw it out :(
disown car;
}
}

contract Inventory {
```

```
 ...
transaction contains(Car@Unowned c) returns bool {
return true;
}

transaction remove(Car@Unowned c) returns Car@Owned {
// Placeholder: implement business logic here
return new Car();
}
}

contract Dealer {

transaction requestToBuy (Car@Unowned car, int price)
returns Order@Shared
{
//assert(car.manufacturer == m);
Order@Shared order = new Order(price, car, this);
car.manufacturer.requestToBuy(order);

return order;
}
}

contract Manufacturer {
Shipper@Shared shipper;
Inventory@Shared inventory;

transaction requestToBuy(Order@Shared o) {
if (inventory.contains(o.car)) {
Car@Owned car = inventory.remove(o.car);
shipper.ship(car);
}
}

}

contract Car {
Manufacturer@Shared manufacturer;
}

contract Order {
int purchasePrice;
Car@Unowned car;
Dealer@Shared dealer;

Order@Owned(int p, Car@Unowned c, Dealer@Shared d) {
purchasePrice = p;
car = c;
dealer = d;
}
}

main contract ... { // Placeholder for main contract }
```

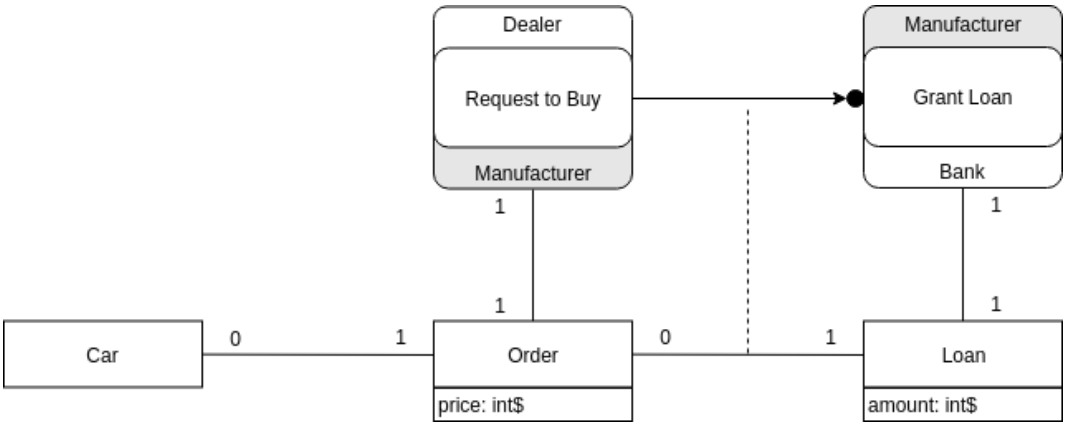

Figure 9 The specification of logging for the loan grant.

4.5.6 An Approach for Specifying Collaboration Logs

The exchange of information mediated by transactions on the blockchain necessitates a guidance towards the information that are going to be written in the ledger, from those that are meant to be kept off-chain. The information provided will serve not only as process documentation but also as a scheme to later extract the necessary information pertaining to process instances, should auditing be required. This becomes of particular significance to enable process mining on data recorded on the blockchain [60, 5].

Figure 9 illustrates an example using a notation inspired by the Object-centric Behavioral Constraints (OCBC) notation [61] and BPMN process choreographies [20]. In the example, we focus on the process fragment pertaining to activities Request to Buy and Loan Money. In the scheme, with $\rightarrow\!\bullet$ we denote that the latter activity cannot take place before the former one, i.e., a *precedence* constraint holds true between records of the corresponding transactions. However, the loan is granted because of a specific sales order, therefore a connection lies not only at a control-flow level, but also as a data-level dependency. Therefore, we specify that the Loan refers to exactly one Order. Likewise, the Order is placed for a Car. The linkage of records in the transactional data will thus require that identifiers of the referred concepts be included in the dependent ones. Additional attributes can be specified along with their data type that are required to be registered in the transactions. In the example, we focus especially on values referring to amounts that the counterparts transfer as value attached to transactions, hence the $ symbol added to the integer data type. In addition to the information pertaining data, transactions record the senders and recipients. We therefore enrich the model by adding them to the activities as in process choreographies. This will clarify the role covered by the accounts involved in the transaction. In the example, the Dealer is the initiator and Manufacturer the recipient of the transaction signifying the Text to Buy activity. Likewise, the bank is the sender of Grant Loan toward the Manufacturer.

4.6 Blockchain Data Analytics: Example of Decentralization of Service-Provider Platform

Alevtina Dubovitskaya (EPFL – Lausanne, CH), Avigdor Gal (Technion – Haifa, IL), Stephan Haarmann (Hasso-Plattner-Institut – Potsdam, DE), Stefanie Rinderle-Ma (Universität Wien, AT), Francesca Zerbato (University of Verona, IT)

Trust and reputation are the pillars, on which online marketplaces are built. Trust is often enabled by reputation [38]. In turn, reputation itself can be defined as *the collection of opinions received from other entities* [25]. Reputation is often used to frame the perception and expectation about someone's behavior based on previous interactions (presumably similar to future ones). Therefore, reputation systems (being an example of collaborative sanctioning/filtering systems) are commonly used to build the trust and to facilitate transactions that happen in online marketplaces.

Many online marketplaces, e.g., eBay, Uber, or Airbnb, use *centralized* reputation management system. In such system, information about the performance of a given participant is collected as ratings from other members in the community who have had direct experience with that participant. The central authority that collects all the ratings then derives a reputation score for every participant, and makes all scores publicly available [30].

In such settings, trust is based not only on the content of the recommendations and rating provided by the platform users. In addition, trust is based on the credibility of the central authority that is solely responsible for ensuring integrity and authenticity of the evaluations provided by the users and for computing the reputation score, and thus enabling the trust. Therefore, an apparent drawback of a marketplace with central reputation-management entity is a single point of failure of the system: can be temporally unavailable or under control of the adversary [7].

To eliminate the risk of having a single point of failure, a central authority that is responsible for reputation management can also be distributed. Indeed, there are environments where a distributed reputation system, i.e. without any centralized functions, is better suited than a centralized system. Distributed reputation systems rely on distributed communication protocol, and reputation computation method used by each individual. However, in case of orthogonal interests of the participants (such as "Hosts" and "Guests" of Airbnb platform), members of the platform can be assumed to be non-trusted, thus, we cannot just rely on them for computing the reputation score correctly. Emerging blockchain technology offers a way to execute processes in a trustworthy manner even in a network without any mutual trust between nodes [40].

Blockchain is a distributed technology that employs cryptographic primitives, and relies on a specific membership mechanism and consensus protocol [12] to maintain a shared, immutable, and transparent append-only register [31, 45]. Data, in the form of digitally signed transactions broadcasted by the participants, are grouped into blocks chronologically and time-stamped. A hash function, applied to the content of the block, forms a unique block identifier, which is stored in the subsequent block. A possible modification of the block content can be easily verified by hashing it again, and comparing it with the identifier from

[7] Due to the nature of the Internet, Airbnb cannot guarantee the continuous and uninterrupted availability and accessibility, https://www.airbnb.com/terms

the subsequent block. The blockchain is replicated and maintained by every participant. A malicious attempt to tamper the information stored in the registry will be noticed by the participants, thus guaranteeing immutability of the ledger. Smart contracts defined to execute arbitrary tasks, enable implementation of desired functionality on top of blockchain.

We can distinguish between *permissionless* and *permissioned* (*public* and *private*) blockchain systems. A system is permissionless when the identities of participants are either pseudonymous or anonymous [57], so that every user may participate in the consensus protocol, and, therefore, append a new block to the ledger. In contrast, in a permissioned blockchain identities of the users and rights to participate in the consensus (writing to the ledger and/or validating the transactions) are controlled by a membership service. A permissioned blockchain is *public* when anyone can read the ledger but only predefined set of users can participate in the consensus, and *private* when even the right to read ledger is controlled by the membership/identity service.

Employing the blockchain technology in the environment of non-trusted or competing entities that rely on reputation management can bring the following benefits. Blockchain by the means of smart contracts can provide transparency and credibility to the way the reputation score is calculated. In addition, the distributed ledger will guarantee immutability and availability of the history of ratings, and reputation scores. However, applying blockchain technology is not straightforward, especially in cases when the data flowing in the system have different levels of sensitivity, volumes, and dynamicity. Moreover, depending on the choice of the blockchain technology implementation, different data-management challenges may need to be addressed.

There already exist an ongoing effort for creating a distributed platform[8] using blockchain technology to provide similar services as the Airbnb platform. Yet, the following question remains: can a decentralized ledger substitute the role Airbnb plays in bringing hosts and guests together? [48] The goal of this work is to analyze different approaches of applying the blockchain technology in the real-word settings, using the example of online marketplace that employs reputation management, i.e., Airbnb. We would like to examine in detail the challenges that arise when permissioned and permissionless blockchain technology implementations are used. Then we propose potential approaches to address the identified challenges, and draw the directions for the future research in the area of blockchain data-management for reputation-based systems.

4.6.1 Ensuring required properties of reputation-management systems using blockchain

In this section, we first provide an overview of Airbnb platform, which is used as an example of an online service-provider platform that relies on centralized reputation management approach. We reason about advantages of decentralizing such service-provider platform. Then, we list the required properties of a reputation-management system, and propose the approaches to ensure them using blockchain.

The Airbnb platform is an online marketplace that enables registered users (“Members”) and certain third parties who offer services (“Hosts”) to publish such Host Services on the Airbnb platform (“Listings”) and to communicate and transact directly with Members that are seeking to book such Host Services (Members using Host Services are “Guests”). Host Services may include the offering of accommodations and a variety of other travel and non-travel related services. [9]

[8] https://www.beetoken.com
[9] https://www.airbnb.com/terms

Airbnb platform, as a service provider, claims no responsibility regarding different kinds of situations that may happen. The Airbnb service can also be temporally unavailable and does not control or guarantee *"the existence, quality, safety, suitability, or legality of any Listings or Host Services, (ii) the truth or accuracy of any Listing descriptions, Ratings, Reviews, or other Member Content..., or (iii) the performance or conduct of any Member or third party"*. In the terms of service, it is stated that the platform is independent, meaning that no endorsements are provided to any of the registered users and the users only may receive some help in facilitation of the resolutions of disputes. [9]

While this centralized and neutral approach facilitates the management of the online marketplace, it can cause issues such as fake listings (and web-pages duplicating genuine Airbnb web-page) created with non-existing properties, problems with local regulations, fake reviews, non-existing users, and non-longevity of the user identity, or listings [33]. The last three issues can sabotage trust derived from the reputation-management system relying on ratings and reviews. For such a centralized platform, it is probably almost impossible to provide the control over and guarantees regarding the behaviour of users, and the quality of the listings provided, due to the highly distributed geographic location of users in countries with different laws and regulations. Decentralization of the platform already could possibly improve the current situation by capturing specifics of the geographic regions, by involving users on a local level to contribute to verification of genuineness of user's identity, listings, and ratings. This will lead to a correct functioning of the reputation-management mechanisms and the increase of transparency and trust.

Resnick et al. [52] claim that reputation systems must have the following properties: *(i)* Entities must be long lived: it should be impossible or difficult for an entity to change identity or pseudonym and erase the connection to its past behaviour; *(ii)* Ratings about current interactions are captured and distributed; *(iii)* Ratings about past interactions must guide decisions about current interactions.

The first property is difficult to achieve in an online marketplace. It is hard to control digital identities unless the binding to the real certificate/ID exists. The latter could be provided after verification using local administration databases, and storing such information in an immutable manner. Second property required availability of the entity/entities that capture and maintain the ratings. While centralized approach may fail due to (temporal) unavailability, in distributed settings the challenge is ensuring consistency of the data among distributed entities. In addition, the second property also depends on the willingness of participants to provide ratings, for which there must be some form of incentive. The third property depends on the usability of reputation system, and how people and systems respond to it [30].

Based on the identified need to improve trust management, and taking into account required properties of reputation systems, we propose to employ the blockchain technology for the following purposes:

- *(i)* verifying existence of the user identities and listings: for example, when booking a property, the user wishes to verify it exists and that the offer is valid;
- *(ii)* validating the quality of the the unit (apartment or service) and its correspondence to the information provided in the listing (reputation of the user can also be used to determine different user roles in the network: more reputable users have more "stake in the network", thus will not be willing to sabotage the system);
- *(iii)* ensuring transparency in communications and credibility of the rating-management system;
- *(iv)* improving the conflict resolution process by, e.g., providing guarantees regarding user payments: smart contract can be used as escrow for safeguarding the payment and setting the terms of payments (the terms can be negotiated).

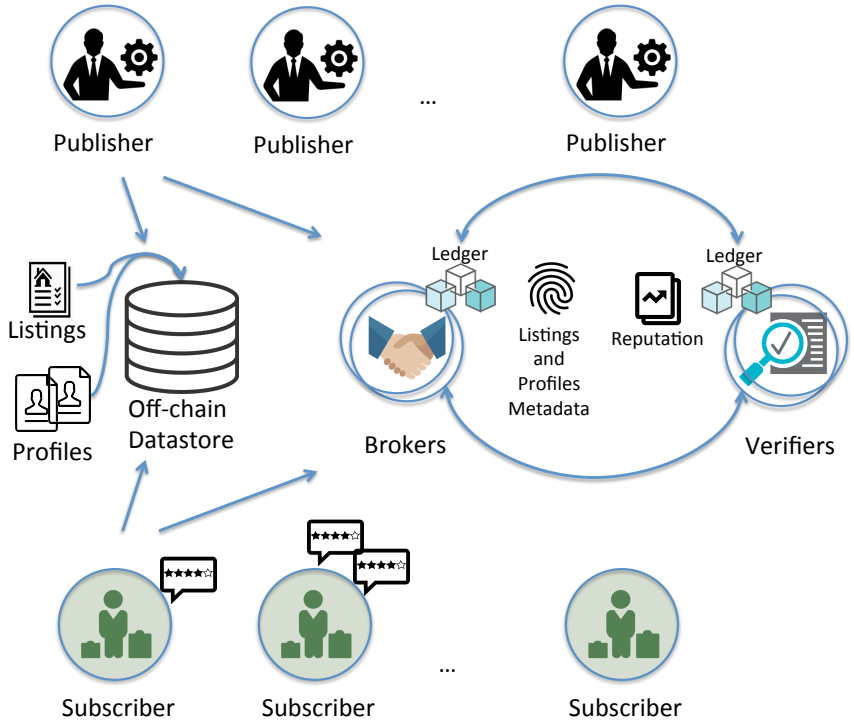

Figure 10 System Model.

4.6.2 Decentralization of a service-provider: system model and design goals

Following the organization of Airbnb service-platform, we consider *Publishers*, with the "Host"-role, and *Subscribers* of the platform, with the "Guests"-role. Publishers are the owners of the property or providers of a service that they want to advertise on the platform, and Subscribers are the renters of the property or the users of the services. In addition, we define the following stakeholders: *Brokers* – entities that may be used as the intermediary between Publishers and Subscribers, and *Verifiers* of the content of the published information ("Listings"), and the correctness of the data (and their correspondence to the physical world). For a user to become a Verifier, a certain level of reputation is required. Verifier can also use external trusted database such as the database of the housing properties registered in the local administration. Publishers and subscribers have orthogonal interests. Hence, neither of them can be fully trusted. Moreover, there is competitiveness within each group: subscribers are looking for the best available listing, and publishers are willing to rent a property or provide a service to the most reliable subscriber.

Next, we list the design goals (DG) regarding the functionality of a blockchain-based distributed service-provider platform. The following actions are required to be available for the users by the platform:

- *DG-1*: registration of a user, ensuring longevity of the user identifier, and preventing from leaving the network and joining with another;
- *DG-2*: registration/removal of a service, or a property to rent out (posting a listing);
- *DG-3*: searching for a service or a unit to be rented;
- *DG-4*: verifying existence of a property of service;

- *DG-5*: booking (negotiation regarding renting conditions such as payment mode, precise timings of arrival, transfer of the key, confirmation, cancellation);
- *DG-6*: payment settlement (with or without involving cryptocurrency, but with respect to the agreement reached during negotiation);
- *DG-7*: providing evaluations and ratings (for publisher by subscriber and vice versa) and computing reputation;
- *DG-8*: detection of the collusion between users and malicious behavior, conflict resolution.

Trying to attain these goals we will rely on the following properties of the blockchain technology. Immutability and append-only properties can be leveraged *(i)* to ensure the history of all the registrations and ratings, and *(ii)* to keep track of the updates of the listings, which can be performed via smart-contract functionality. Smart contracts will be used as well for transparent computation of reputation (based on the provided ratings), for negotiation of the terms of booking, and settling cancellation. Potentially, the full history of renting, ratings, and communications between users (for instance, when negotiating booking conditions) can be leveraged using machine learning approaches to assist subscribers searching for a unit or a service and to detect the malicious behavior trends.

4.6.3 Challenges when applying blockchain

While the advantages of employing the blockchain technology listed above advocate for its adoption, in order to achieve the design goals, the following challenges have to be addressed first.

- *C-1: Choice of the blockchain-technology implementation.* The choice of technology highly depends on the use-case scenarios. It is crucial to define the mapping between the "peers" and the real-world entities and to define who will maintain the blockchain, i.e., who are the peers/entities that will be storing a distributed ledger? Depending on the sensitivity level of the data that flow in the system, as well as the degree of the involvement of the peers, next step is setting up the the policy to join the network and to read/write the transactions from/to the ledger.
- *C-2: On- and off-chain data management.* It is highly important to define what kinds of data will be stored on- and off- blockchain to avoid unnecessary data replication that can make a system impossible to use in real-world settings due to high latency, privacy issues, and requirements to process big volumes of data. Definition of the structure and formats of on- and off-chain data depends on how the following challenges are addressed:
  - *C-2a: Translating business processes and negotiation terms to smart contracts.* Desired functionality of the system and capabilities of the smart contracts to capture it have to be considered first, to understand what is the minimum required amount of the data that have to be stored on the distributed ledger. For instance, in the case of Airbnb-like platform, we assume a possibility for negotiation of booking conditions. Therefore, the terms to be negotiated have to be taken into account beforehand. Using smart contracts for automatic payment settlement requires financial data about the user to be available in the system and coordination with banks and companies providing online payment services.
  - *C-2b: Ledger design and data expressiveness.* The ledger is the list of append-only transaction, together with read-write set that is replicated among all the entities that maintain the ledger. Therefore, transactions organized in such a way are not efficient to query. The design of the read-write set (the data stored on the blockchain) has to ensure efficient queries and required functionality by defining the database structure

and data organization accordingly. Storing the whole data lake with all the listings provided by users and all the communications between Publishers and Subscribers is impractical, while keeping track of the users registration, or having metadata, ratings and evaluations is feasible.

- *C-2c: Data privacy and security.* The data stored on-blockchain and replicated among distributed entities can have different sensitivity levels. During user registration, as well as verification of user's identity, or genuineness of the property or a service described in the listing, access to the highly sensitive data are required. The listing itself, once its validity is ensured, may contain only publicly-available data. Applying cryptographic techniques can become necessary to enforce access control policy and ensure data privacy and security. This, however, introduces the challenge of managing cryptographic keys during their whole life-cycle.

- *C-3: Consistency between digital representation of objects and the real-world.* Establishing and maintaining consistency between digital and physical worlds is challenging. Information about all the registered users/units/services has to be verified using some trusted sources, and the information about registered users and their reputation. With the absence of a centralized entity, it is also challenging to maintain consistency and handle conflict resolution.

The fundamental and first-to-address challenge is defining the choice of the technology and setting up the network and policies. Depending on the concrete use-case scenario and how *C-1* challenge is addressed, i.e., what type of blockchain technology implementation is chosen, different considerations have to be taken into account when resolving challenges *C-2* and *C-3*. We next discuss how the choice of blockchain technology influences challenges *C-2* and *C-3* and the potential approaches to address them.

4.6.4 How to address data-management challenges when implementing distributed blockchain-based service-provider

The choice of technology shapes further challenges related to the blockchain data management. First, we provide more details about differences between permissioned and permissionless blockchain technology implementations. Then, we discuss research directions and potential approaches to address the aforementioned challenges *C-2* and *C-3* in the framework of applying permissioned and permissionless blockchain technology.

Blockchain technology implementations: Below we briefly describe the characteristics of the permissionless and permissioned blockchain technologies using their existing implementations with smart contract functionality as an example.

Ethereum[11] is an implementation of a permissionless programmable blockchain that enables any user to create and execute the code of arbitrary algorithmic complexity on the Ethereum platform: Ethereum Virtual Machine (EVM). EVM can be seen as a large decentralized computer. "Accounts" of two types could be created on EVM. Externally owned account (EOA) is an account controlled only by a private key of a user. The owner of the private key associated with the EOA can remain anonymous (up to a certain degree) and has the ability to send messages. Contract account is the second type of accounts that can be seen as an autonomous agent that lives in the Ethereum execution environment and is controlled by its contract code: smart contract. Smart contract is used to encode arbitrary state transition functions, allowing users to create systems with different functionalities by transforming the logic of the system into the code. In case of public blockchain (such as Ethereum Mainnet), smart contracts and all the transactions are public.

In Ethereum, transaction processing is Turing-complete and it can be used to implement any public functionality in a distributed fashion, but the code execution must be paid. The transaction price limits the number of computational steps for the code execution in order to prevent infinite loops or other computational wastage. Users can participate in the consensus process to obtain the tokens in order to pay for the transaction execution. In Ethereum, the consensus is achieved by using GHOST – modified proof-of-work (PoW) mechanism.[10]

In order to avoid issues of network abuse, all programmable computations in Ethereum are subject to fees. The fee schedule is specified in units of gas. Thus any given fragment of programmable computation (i.e., creating contracts, making message calls, utilizing and accessing account storage, and executing operations on the virtual machine) has a universally agreed cost in terms of gas [66]. Ethereum provides excellent scalability in terms of number of nodes and clients, but has a limited transaction throughput. In 2016, typical Ethereum throughput was fewer than 20,000 transactions per day, i.e., about 0.2 tx/s on average [63].

In contrast, due to the architectural design and different type of consensus protocols employed, scalability in terms of number of nodes in case of permissioned blockchian technology is limited. *Hyperledger Fabric* [24, 2] – an implementation of a permissioned blockchain – is an open source blockchain initiative hosted by the Linux Foundation. Hyperledger Fabric contains a security infrastructure for authentication and authorization (membership service). It supports enrollment and transaction authorization of peers and users through public-key certificates. This is one of the main differences with the permissionless blockchain framework. In Hyperledger Fabric, in addition to the membership service, the other main architectural components are peers, and ordering-service node, or orderer. Orderer is a node running the communication service that implements a delivery guarantee, such as atomic or total order broadcast. The ordering service can be implemented in different ways, ranging from a centralized service to distributed protocols that target different network and node fault models.

Architecture design of permissioned blockchain technology may introduce a certain level of centralization due to relying on the membership service and orderer (that, can also be distributed to prevent a single point of failure in the system). However, such a design provides privacy and security guarantees that are impossible to achieve in the permissionless settings.

Smart contracts are implemented by the chaincode that consist of Logic and associated World state (State). Logic of the chaincode is a set of rules that define how the transactions will be executed and how the State will change. The Logic can be written using general-purpose programming language. The State is a database that stores the information in a form of key-value pairs, where the value is an arbitrary byte array. The State also contains the block number to which it corresponds. The ledger manages the blockchain by including an efficiently cryptographic hash of the State when appending a block. This enables efficient synchronization if a node was temporary off-line, minimizing the amount of stored data at the node.

On- and off-chain data management and consistency between the ledger and the real world: In both settings, given the data volumes managed by the service-provider platform in an online marketplace, it is impossible to guarantee data privacy, consistency and efficient queries over the data, while keeping all the data on-chain. Innovative approaches for storing, indexing and describing the data are thus required to ensure privacy, consistency and efficiency. However, we could also leverage the possibility to store the data off-chain and the

[10] http://www.ethdocs.org

design and structure of the ledger to achieve aforementioned properties. Next, we discuss how could we address the challenges of on- and off-chain data management and achieving consistency between the ledger and the real-world within permissionless and permissioned settings.

Permissionless settings. Based on the characteristics and properties of the permissionless blockchain, we would like to point out several research directions and potential approached that could be applied regarding on-/off- chain data-management challenge:

- Everyone can join the network, as there is no centralized entity, the mechanisms embedded in the smart contracts have to be developed to make sure that users identifiers are static. Pseudonymization approach could be used: such as creating multiple pseudonyms to ensure privacy, but that can be linked when computing reputation.
- Currently, PoW is used: transaction fees are involved. Therefore, we have to take this into account when writing smart contracts, as every operation have to be paid.
- Sensitive data, such as exact locations and names of the guests, should not be placed in plain sight on the ledger due to privacy issues. Therefore, there is a challenge in finding a way to perform operations over encrypted data (lightweight homomorphic encryption), or over statistical (aggregated anonymized) data.

Premissioned settings. For the case of permissioned settings we make the following considerations:

- The policy of joining a network and differentiating users' rights into those who can add new blocks, submit a transactions, read the ledger, *etc.*. A possible approach is to let only users with the certain level of reputation to maintain the ledger. However, incentive mechanisms and rewards have to be defined.
- In general, in case of permissioned blockchain less peers will be maintaining the ledger (data replication is simplified, however, still depends on the consensus mechanism employed). All of them will also be registered. Therefore, access control policy can be partially enforced already by membership service. Implementation of the cryptographic techniques in the distributed environment with less peers is more practical, yet further investigation are necessary.
- Membership service could also be employed to ensure longevity of the user's identifier. Yet the exact requirements for constructing pseudonyms/identifiers have to be developed.
- Digitization of business processes (e.g.,verification and publish the listing) can be done via smart contracts and secured with endorsement policy (such as in the case of consensus protocol employed in Hyperledger Fabric v1.X [2]). For instance, verification of a listing will require the endorsements from *(i)* trusted sources, such as government registry and *(ii)* a number of local peers with a certain level of reputation.

Consistency between the ledger and the real world: The *C-3* challenge related to ensuring that the ledger stores genuine information is very hard to achieve, in both settings. Smart contract functionality can be employed to automatize the process of verification of the information using trusted databases, or the local peers with a certain level of reputation. For such verification to be reliable, the system has to ensure collusion detection and provide a possibility to exclude the malicious/colluding peers. In permissioned settings, for instance, this can be achieved through the policy defined on the side of the membership service.

4.6.5 Conclusion

The expansion of blockchain-based applications proliferates in different areas, and online marketplaces are not an exception. Centralized service providers, especially the ones relying on reputation and trust, will definitely benefit from adopting the principles of the blockchain technology, and the properties it brings, namely transparency, immutability, and credibility, to name a few. Yet, multiple data management challenges arise when such decentralization occurs.

In this work, we analyzed these challenges using a practical potential use-case scenario of Airbnb service-provider platform. We proposed possible ways to address these challenges in different permissionless and permissioned settings, indicating the directions for the future research in the area of blockchain data-management for reputation-based systems.

Addressing the challenges only from the data-management perspective may not be enough. However, it could simplify compliance of such applications whit the local laws and regulations. Intelligent data management, together with transparency and credibility brought by the blockchain, opens the doors for leveraging machine learning techniques to enhance and facilitate the use of the online marketplaces, for instance to enable recommendations of the unit or the service that can be of interest for a subscriber.

4.7 Data Technology to the Rescue: Digging the D in GDPR

Søren Debois (IT University of Copenhagen, DK), Alevtina Dubovitskaya (EPFL – Lausanne, CH), Avigdor Gal (Technion – Haifa, IL), Petr Novotny (IBM TJ Watson Research Center –Yorktown Heights, US), Stefanie Rinderle-Ma (Universität Wien, AT), Stefan Schulte (TU Wien, AT), Ludwig Stage (Tübingen, DE), Kaiwen Zhang (ETS – Montreal, CA)

The motivation for the discussion presented in this section is the recently introduced GDPR (General Data Protection Regulation) 2016/697, an EU regulation that aims at regulating the way personally identified data is being gathered and consumed and to define the legal rights of people to the use of their data. Taking the point of view of computer and data scientists, we wish to identify suitable mechanisms to support organizations in their journey towards the compliance of their information systems with GDPR.

The discussion involved a deeper understanding of the GDPR and the requirement it puts with respect to data protection. Equipped with this understanding, we have identified elements of matured (databases) and new (blockchain) technologies that can be put into use when adapting an organization's information system to be GDPR compliant.

4.7.1 GDPR 101

The EU General Data Protection Regulation [1] came into force May 25, 2018. It applies within the EU and the European Economic Area; however, because information-centric businesses tend to be global, it is of global concern.

The GDPR confers onto citizens (*data subjects*) a number of rights, and onto companies and institutions using that data (*data controllers*) a number of obligations, the latter at the penalty of potentially significant fines.

Of particular interest to this chapter are the following stipulations of the GDPR. It is worth noting that our discussion is not comprehensive.

- GDPR is concerned exclusively with personally identifiable data. It specifically states that anonymisation stops data from being personally identifiable. Pseudonymisation does not, but is considered, in some cases, an adequate security measure [43].
- Any processing of data must be for a specific purpose, and that purpose must be legitimate. Common legitimate purposes include those that are necessary towards fulfillment of a contract (*e.g.*, I must record your address to ship you goods) as well as those required by law, those required for ongoing court cases, *etc.*
- Data subjects have a number of rights:
  - The right to erasure (Article 17): When a purpose ceases to be legitimate, perhaps because it has run its course, or perhaps by request from the data subject, a data controller must erase **without undue delay** personal data.
  - The right to rectification (Article 16): When given notice of incorrect data, a controller must, again **without undue delay** rectify incorrect data.
  - The right to data portability (Article 20): Upon request, a data subject may receive from a controller all of hisqhers personal information in electronic form.
- GDPR imposes a number of obligations on data controllers. The group discussed the following:
  - The obligation to keep records of processing activities (Article 30).
  - The obligation to inform data subjects about processing (Article 13, 14)
  - The distinction between the data controller and its data processors, and the requirement that a data controller has a contractual agreement with its processors.

Controllers and Data Processors

The GDPR poses many challenges on interoperability between partners exchanging data. Partners can basically have two roles as defined in Article 4 of the GDPR: a controller "which [...] determines the purposes and means of the processing of personal data" [11] and a processor that "processes personal data on behalf of the controller". In addition, there might be a supervisory authority that monitors the compliance with GDPR in an EU member state. The controller and the processor(s) might have to cooperate with the supervisory authority. It is also possible that several controllers jointly determine the purpose of the processing, then they are called joint controllers. In agreement with the (joint) controller the processors can engage other processors.

Figure 11 shows a potential network of partners. All engagements in such a network *e.g.*, between controller and processor) are subject to legally binding contracts with respect to the GDPR and thus in essence with respect to the use of the data. Article 30 states that the controller "shall maintain a record of processing activities under its responsibility" and the processor "shall maintain a record of all categories of processing activities carried out on behalf of a controller". Already these two basic obligations (establishing contracts and logging) for data exchange under GDPR point to the usage of blockchain technology due to its support for decentralization, trust, and transparency. Firstly, the legally binding contracts between controllers and processors, and possibly also between joint controllers and processors that engage other processors, could be possibly implemented as smart contracts. Secondly, the corresponding logging of the different partners (with different roles) could be realized using blockchain technology. Despite the aforementioned potential benefits, details of the realization and the resulting complexity are yet to be investigated.

[11] https://advisera.com/eugdpracademy/gdpr/definitions/

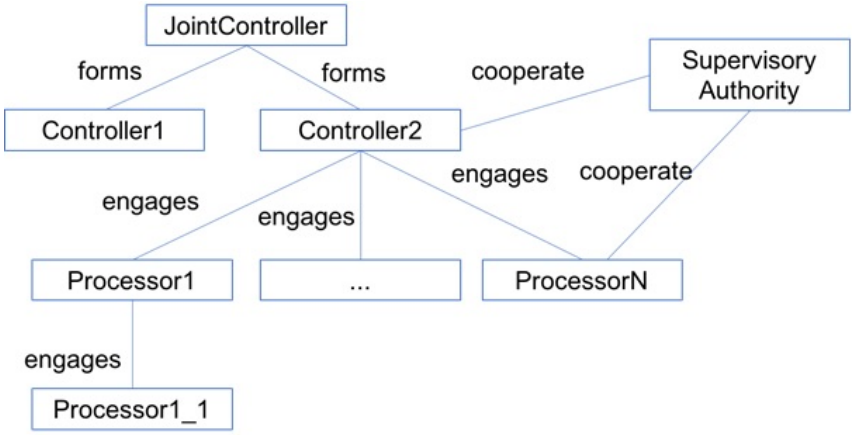

Figure 11 Partner interoperability

4.7.2 Technological Background: Data Life Cycle, Databases, and Blockchains

The data science discipline defines a lifecycle of data that captures roughly four steps: data is first **gathered** using possibly sensors, human input, or otherwise already available in data stores (such as the Web). The gathered data is **managed**, integrated, and stored on a facility that may range from a personal storage to a cloud storage, by means of data management systems such as a database management system (DBMS – see below). The stored data is prepared to be **analyzed** by machine learning algorithms and the outcome is finally **presented** to users by means of database queries, managerial dashboards, *etc.*

Database technology has been around for many years now (*e.g.*, [51]). Using a DBMS, one can define a *schema*, describing the content of the database using a conceptual model such as the relational model. Data is accessed via *queries*, which can either be retrieval queries or update queries (insertion, modification, and deletion). A *log* of all update activities is kept to assist in situations of crash recovery or other failures. Finally, data may be distributed over multiple sites and partially replicated to ensure speedy retrieval and to guard against failures.

The emergence of blockchain technology has opened manifold opportunities to redesign collaboration. In general, blockchains allow to store data in a distributed way, with each participant in a blockchain network being able to possess the complete blockchain (allowing transparent data sourcing) and verify the stored data. Blockchain technology does not rely on trusted third-party architecture, which creates entry barriers and a single point of failure. Rather, blockchains guarantee the integrity of the data and smart contracts enable distributed execution without delegating trust to central authorities nor requiring mutual trust between each pair of parties. Furthermore, blockchain technology potentially enables fine-grained access controls, allowing different parties to selectively share different data with different partners, using data that is resident in the blockchain.

Three specific elements of blockchain are important to the following discussion. First, blockchain can serve as a trusted (ledger) data storage. No data, once written to the blockchain, can be modified without letting everybody with access to the blockchain knowing about it. Second, the use of smart contracts, computer scripts that define a sequence of reactive operations on which partners agree ahead of time, enable a transparent operation of

a system, subject to the audit of all participants. Finally, blockchain employs mechanisms for replication of the data on a blockchain that allow the continuity of recording activities even when faced with network instability.

There is a common distinction between *permissionless* and *permissioned* (*public* and *private*) blockchain systems. A system is permissionless when the identities of participants are either pseudonymous or anonymous [57], so that every user may participate in the consensus protocol, and, therefore, append a new block to the ledger. In contrast, in a permissioned blockchain identities of the users and rights to participate in the consensus (writing to the ledger and/or validating the transactions) are controlled by a membership service. A permissioned blockchain is *public* when anyone can read the ledger but only predefined set of users can participate in the consensus, and *private* when even the right to read ledger is controlled by the membership/identity service. In this work, we refrain from discussing the various design options that come with blockchain system being premissionless or premissioned.

In public networks like Bitcoin and Ethereum, storage of large amounts of data is expensive. Therefore, there is a need to rely on off-chain file sharing networks, such as IPFS [7], Filecoin,[12] and Swarm.[13] The basic functionality is to record a hash of the document on the blockchain, send the original document on the file sharing network, and retrieve the document using the content address which was previously stored on the blockchain.

In IPFS, there is an assumption that participating nodes are altruistic, and are willing to store data simply because they wish to maintain the network available. However, this altruistic model is weakly applicable in practice, with files usually remaining available for 24 hours at most. To obtain longer availability, users rely on pinning services, which will actively maintain the information long-term, in exchange for financial compensation.

To address these issues, newer systems like FileCoin and Swarm have built-in incentive mechanisms to promote long-term file availability.[14] In FileCoin, participating nodes employ Proof-of-Storage to mine blocks and collect mining rewards. Mining power is calculated based on the amount of data stored, which can be verified using Proof-of-Retrievability. In Swarm, positive incentivization is provided when serving content, which incentivizes nodes to retain popular files. In addition, negative incentivization is enforced through a staking (security deposit) mechanism, where nodes have to commit some cryptocurrency resources in order to participate. An audit mechanism then allows the data owner to challenge storage nodes, who must provide the requested data, or have its stake slashed.

4.7.3 Personally Identified Data

Putting personally identifiable information on a public blockchain in an immediately readable (*i.e.*, non-encrypted) form is likely prohibited specifically because (i) a blockchain, as an immutable database, can never satisfy the right to erasure, and (ii) because it would be impractical for a data controller to obtain a contract with every node in the blockchain as a data processor. These limitations may be resolved when the data are stored and or modified such that the personal identifying information is not present. Moreover, the willingness of individuals to agree to the data collection can be increased with such a solution.

[12] https://filecoin.io/
[13] https://swarm-guide.readthedocs.io/en/latest/introduction.html
[14] http://swarm-gateways.net/bzz:/theswarm.eth/ethersphere/orange-papers/1/sw%5E3.pdf

Tokenization is a technique, aimed at separating the true identity of a person and its representation in a database.[15] Tokenization has two main variations, namely anonymization and pseudonymization. The use of tokenization mechanisms allows to disconnect the various types of data about a person (such as income, education, *etc.*) from the data uniquely identifying the person (such as social security number). The mapping information between the various types of data is separated from the personal identifying data and stored within a secure storage of the tokenization system. This approach allows the use of personal data in data processing tasks without revealing the actual identity.

Tokenization may allow re-identification [55] of personal data from tokenized data, especially if the same pseudonym is used repeatedly for the same person. Successful re-identification does not expose personal data directly since there is no personal data stored on the blockchain. What would be leaked is information about what kind of data is stored for a particular person and for which purposes, which may, by itself create a breach of GDPR regulations. To provide controlled linkability, Camenisch and Lehmann proposed a combined approach of pseudonymization with a potentially untrusted server that stores the mapping [13]. Pseudonyms should be unlinkable by default, yet preserve the correlation that enables to re-establish the linkage only if necessary and based on the policy specified via a (potentially untrusted) converter. The converter establishes individual pseudonyms for each server derived from a unique main identifier that every user has, but without learning the derived pseudonyms. The converter is still the only authority that can link different pseudonyms together, but it does not learn the particular user or pseudonym for which such a translation is requested. To construct such framework, the authors use dual-mode signatures (which allow one to sign messages in the plain as well as when they are contained in an encryption), (verifiable) pseudorandom functions, and homomorphic encryption.

In case of a unique mapping, all historical activity can be traced, which is both advantageous and undesirable at times, for example when it comes to preserving privacy. In medical studies, double-blind studies require that even the medical personnel should not know who is actually provided with a novel medical treatment and who is part of the control group. AnonRep [69] is the first practical anonymous reputation system maintaining the unlinkability and anonymity of users' historical activities. AnonRep uses verifiable shuffles and linkable ring signatures, with a multi-provider deployment architecture.

Similar to the right of data erasure, data subjects can request that only anonymized data can be used for the data analysis tasks. This is also somehow beneficial for the data controllers: if the data are anonymized, there is no need to comply with GDPR.

How can we actually guarantee that the data are properly anonymized and re-identification of the data subject is impossible, taking into account that more data about the same data subject can be revealed in the future? Depending on the nature and sensitivity of the data, and the field of the research question for which the data are used, there are various considerations such as what is a required privacy level and how to choose the parameters of the anonymization algorithms to achieve this level of data anonymization (*e.g.*, k-anonymity). Blockchain technology can be used in order to track the data that have been released in the anonymized form, including the anonymity level of the data (but not necessary keeping all the data, only some metadata). This will be used as an input to compute the risk of being re-identified if more data about the data subject are released. Kiyomoto *et al.* [32] propose a blockchain-based distribution scheme for anonymized datasets. The platform consists of peers that act as data brokers, data receivers, and verifiers of transactions, and

[15] https://en.wikipedia.org/wiki/Tokenization_(data_security)

blockchain is used for recording all transactions of anonymized datasets between a data broker to a data receiver. One can argue that keeping track of all the transaction can put the anonymization at risk. However, as mentioned above, recording of all the transactions can be used to compute the risk and prevent the violation of the data subject's privacy.

Storing access control policy/permissions and transaction history on the blockchain has the benefit of letting the user transparently identify who has/had access to her data and for which purpose. Sensitive data may also be stored on blockchain, but in this case the data must be encrypted, key management and access control mechanisms must be put in place (for instance, via smart contracts).

4.7.4 Data Gathering and Usage

Within the GDPR, data can be processed, including being gathered, only for a legitimate purpose, to be clearly defined in advance. Given a purpose and the necessary permissions, a DBMS may be designed, creating a database schema according to which data is collected using database insertion and update queries.

Recall that a database log details the sequence of activities (insertions, updates, and deletions) performed by the database. The append-only character and the property of immutability of blockchains make them suitable for logging purposes. Logging only needs to append data, and logging an update is simply done by creating a new log entry. The immutability feature might be beneficial in terms of creating a tamper-free audit trail that could establish a higher level trustworthiness in an audit report done with the help of such a log.

A log responsibility can be extended to serve as a mechanism for ensuring that data is used solely for its legitimate purpose by using a blockchain to store the log. The extended log (performed using an API/gateway/wrapper) will include all the queries that were performed throughout history on the database. Its append-only nature, combined with mechanisms to avoid tampering with the registered data, allows a trustworthy mechanism to record database queries for possibly future auditing and other tasks that are relevant to GDPR.

When using blockchain to store a DBMS log in a trustworthy fashion, three main questions come to mind. First, how can one trust the controller or any of its contractual data processors to faithfully use this mechanism? This is, in fact, a question that is beyond the scope of GDPR. We note here that GDPR is meant to define the procedures that are needed to be set in place to ensure the privacy of personally identifiable data. The enforcement of such mechanisms, once defined to be in place, is left to be performed by other conventional means already specified by law.

Second, how can one guarantee the correctness of the log in the face of database failures. For this, a rich literature exists, on the use of a *write ahead log* to allow safe recording of database activities.

Third, how can one ensure the correct recording when it comes to blockchain underlying mechanisms. In particular, how can one ensure that blocks containing relevant log information will not be discarded. This can be possibly done by using an incentive mechanism, inherent to blockchain, which ensures a positive payment for storage.

GDPR internal and external auditing can be done using the blockchain-based storage and the software systems, thus certifying that a company has put the technical means in place in order to comply with GDPR requirements.

4.7.5 Right to Erasure

The GDPR provisions for the erasure of personal data upon request. Any mechanism that supports the right to erasure should also be able to provide a positive proof to the erasure that will serve for auditing purposes.

Data management involves a set of steps that are aimed to ensure the storage of a semantically meaningful data. This is typically done by the use of a DBMS, which capabilities provide semantic guarantees over the data. One such guarantee is the maintenance of integrity constraints, allowing to maintain a connection among various elements of data. When it comes to the Right to Erasure, such a request is translated into a delete query in the database. The query is duly reported in the log (now stored on blockchain), which can be served as a proof of compliance with the request. A point to keep in mind is that whenever required by law (*e.g.*, for tax purposes or money laundering regulations) data cannot be erased in response to a request for erasure. In such a case, the database consistency may be infringed. Such situations can be handled by advanced techniques for exception handling in database consistency.

Data in databases may be distributed (using distributed database techniques) and partially replicated on a peer system (using methods that are similar to IPFS). Such an architecture requires specialized mechanisms to deal with erasures on multiple peers. This can be possibly done by using an incentive mechanism, where data that is intended for erasure will no longer receive positive payment for storage and will eventually be dropped from the peer network. We note that such a mechanism is in line with GDPR, which requires that the request for erasure will be performed without undue delay, rather than immediately.

The use of blockchain technology for storage of personal data thus must consider the implications of storing data onto an immutable ledger and employ appropriate mechanisms to allow the erasure of personal data. Blockchain technology supports solely an append-only mechanism for storing data. This presents a challenge when dealing with the right to erasure, since all queries are recorded in the log. The use of sophisticated tokenization can assist in upholding the right to erasure. When using anonymization, the identification of a person cannot be recovered. When using pseudonymization, simply erasing the link between the pseudonym and the identifier can do the trick and turn the pseudonym into an anonymized entity, no longer traceable to the person it represents. In addition, there are several known approaches of effective erasure of data stored on the ledger that uphold the integrity of the ledger.

A related topic to the right to erasure is the enabling of an opt-out option from data gathering, even if the data is anonymized. If the data subject agrees to provide anonymized data, but would like to be informed about the use of the data (*e.g.*, the outcomes of a research study), the combination of pseudonymisation and anonymization can be used. If the data about the data subject were used as an input to an algorithm, then a data subject can require to stop processing the data. In such a case, a careful analysis is needed to check a possible impact on the privacy of other data subjects, whose data are processed by the same algorithm and stored in the same database.

Erasure mechanisms for tokens

Proof-of-Burn is a consensus mechanism for Proof-of-Work and Proof-of-Stake, proposed for public cryptocurrency blockchain systems [49]. In this mechanism, coins are burnt by sending them to an unredeemable output. In Bitcoin, this is accomplished by specifying

an output script that will never evaluate to true (*e.g.*, push 4, check if it is equal to 5[16]). Once coins are burnt, a consensus algorithm that functions based on Proof-of-Burn relies on miners to supply the proof that sufficient coins have been burnt (*i.e.*, exceeds the required difficulty) to convince other miners to accept the block as valid. Because coins burning is sent as a regular transaction, it is necessary for the burning transaction to be stabilized (*i.e.*, accumulate sufficient confirmations). Hence, the consensus protocol would set a lower limit on the amount of confirmations necessary for valid proofs.

Proof-of-Burn can be used in the context of a GDPR-compliant public blockchain to support the right to erasure in the context of data tokenization. Given a blockchain platform where tokens have to be generated and redeemed for each piece of data in order to process this data further, a user who requests the right to erasure can notify the network to burn tokens associated to her account (per ID). Each account that currently owns tokens associated with this user must then submit a transaction to burn the tokens. After a stabilization period (measured in confirmations), the user can challenge any participant to demonstrate that her right has been executed. Challenged participants must then supply the appropriate Proof-of-Burn (*e.g.*, address to a transaction who spends the tokens, and sends them to an unredeemable output). Thus, the information embedded in the tokens are no longer usable, which stops further processing of this information.

This approach works insofar that each token output has a finite consumption limit attached to them (*e.g.*, transaction outputs for cryptocurrencies are single-use). For certain type of tokens, it is possible that their consumption is unlimited (the same transaction output can be used multiple times). To accommodate this case, the input script provided by the burning transaction should include a special operation (to be provided by the blockchain platform) that consumes the token permanently, barring future transactions from redeeming it.

One possible issue is that the burning transactions are never accepted into the blockchain, possibly due to low transaction fees attached, or blacklisting policies from certain miners that preclude the transactions from being included. Because unconfirmed transactions will not prevent the tokens from being processed further, it is the responsibility of the data controllers to ensure that the transactions are confirmed, whether this means that the transaction fees have to be raised, or that there is a sufficient proportion of miners who are willing to accept the transactions so that they will eventually be included in a mined block. This could be incentivized by adding special mining rewards, embedded in the blockchain core protocol, for blocks that contained burned tokens. This type of reward, coupled with the special burning operation detailed in the previous paragraph, could be used to demonstrate that a specific blockchain platform is GDPR-compliant.

Proof-of-Burn can also be useful for interoperability between multiple blockchains. In order to transfer a token from one blockchain to another, it must first be burned in the original blockchain. Once this burn transaction has been confirmed and stabilized, a new transaction on the new chain can be used to redeem the same data token, by supplying the Proof-of-Burn referencing the previous transaction. This ensures that the same token is not "double-spent" on multiple blockchains. This concept is already used for cryptocurrencies, when transitioning after a hard fork.

We note here that UltraNote provides self-destructing data storage (based on an expiry time) [59]. However, it is unclear if the underlying cryptomechanisms are usable when erasure is explicitly requested at an arbitrary time.

[16] https://en.bitcoin.it/wiki/Proof_of_burn

Cryptoeconomic erasure for off-chain storage

With a cryptographic approach based on either symmetrical or asymmetrical cyphers, the personal data are encrypted prior to storage onto the ledger. Later, when the data are retrieved from the ledger, the data must be first decrypted to reconstruct the contained information. The data become effectively erased when the decryption key(s) are not available. This approach requires an encryption and decryption key management mechanism, which upon request from the data owner or data manager erases decryption keys and thus erases the data encrypted with these keys. The use of homomorphic encryption allows a limited number of operations on already erased data while providing compliance with the GDPR requirements. The key management mechanism must allow encrypting the various combinations of data in logical units such that the later removal of the decryption keys leads to erasure of the intended (and only intended) data. Moreover, the key management mechanism must be trusted by all participants of the blockchain network and thus appropriate architectural and operational guarantees and policies must be put in place. For example, to provide an independent key management mechanism in a private blockchain network of equal participants, the participants may agree and elect an independent and trusted third party (centralized or distributed) responsible for the key management. It is worth noting that the cryptographic approach, if implemented correctly, provides guarantee of complete data erasure while at the same time requires computationally complex decryption before data can be used in processing, which may be prohibitive in large datasets.

Using tokenization, when data are requested to be erased, removing the mapping information from the tokenization system (*i.e.*, data allowing to establish the connection between the personal identifying data and other types of personal data) causes the data to be effectively erased. In comparison to the cryptographic approach, which requires complex decryption before data can be used, tokenization does not increases the computational complexity of data processing. However, since tokenization leaves the data accessible after erasure of mapping links, it thus requires careful design of the data structures in order to eliminate all connections between the personal and identifying information.

Incentivisation

Positive and negative incentivization mechanisms can be used to promote eventual (long-term) unavailability of specified content. A soft negative incentivization, which does not require any changes to the underlying file sharing network, works as follows. When an erasure is requested, it is stored as a transaction on the public chain, thereby accessible by all storage nodes. At that moment, sharing of the erased content is forbidden: any further sharing of the content can be included as part of an "evidence transaction", recorded on the main chain [8]. A successful evidence transaction leads to the stake of the perpetrator to be slashed. If the penalties are higher than the incentives collected for sharing data, storage nodes will not be motivated to serve "erased" data. Consequently, storage nodes can only store private copies or disseminate it out-of-band. However, if the primary purpose of the storage node is to dedicate its storage resources to the file sharing network, it will eventually evict the data from its system and store newer information that can be monetized. From a public standpoint, the data is effectively forgotten from the system in a probabilistic manner, depending on various factors, such as the rate of incoming data, the average amount of storage resources per node, the average stake deposited per node, and the penalty for serving erased data.

A positive incentivization mechanism requires modifications to off-chain file sharing networks. We demonstrate the modifications using Swarm. An erasure transaction on the main chain will record a bogus hash value associated to the content address. When a Swarm storage node is serving a piece of data, the content hash is verified against the main ledger. If the hashes do not match, the reward for serving content cannot be collected. Thus, storage nodes are incentivized not to store erased data, since they can no longer collect the reward associated with serving them. Furthermore, main chain peers will eventually prune the original transaction (containing the content address) from the Merkle tree of that block, since that address is no longer available or monetizable from the underlying file sharing network. New nodes doing a full sync on the blockchain will simply received a pruned version of the blockchain without the erased data. Old nodes can still retain the erased content addresses, but will eventually forget them during garbage collection processes.

The incentivization mechanisms provide cryptoeconomic incentives to forgetting erased data. However, irrational agents may still choose to keep the data, so it can only be considered probabilistic. We believe such a mechanism satisfies the GDPR requirement of erasure without undue delay.

4.7.6 Right to Data Access and Rectification

Currently, information about personally identified data needs to be requested from the company's individual employees and administrators, which is a very time-consuming (and therefore costly) and also a naturally error-prone approach. The use of a DBMS and its extended log recording on a blockchain would allow an easier response for a data access request. People who are interested in getting information about how their personal data is used within an organization under GDPR's right to data access could simply query the blockchain-based data storage, thus getting this information. The query could be both implemented off-chain as well as be provided as a smart contract. In the latter case, transparency would be given and the user does not have to implement a querying mechanism by herself. However, this approach requires all personal data to be traceable, *e.g.*, by using the same personal ID for all data related to a particular person.

The positive incentivization mechanism can also be used for the right of rectification, by replacing an existing piece of data with a newer one, which will now be monetized.

4.7.7 Conclusion

This work provides a technical support to an effective implementation of the GDPR regulation, using techniques that are based on both blockchain and databases technologies. The paper presents GDPR, the relevant technologies, and the way we envision these technologies can serve an organization in becoming GDPR-compliant.

In our proposed solution we took advantage of the main benefits of each of the two named technologies. From databases, we take their ability to maintain data consistency, as well as its ability to efficiently store and retrieve large amounts of data. Blockchain's ability to secure transactions and keep a trustable ledger complement the set of needed abilities.

There are many open questions left to be discussed in future research. Especially, when realizing the proposed technology many issues may rise. For example, how to establish smart contracts between partners of different roles in exchanging data under GDPR? Also, how to realize (distributed) logging between a) controllers and processors, b) processors and processors, c) joint controllers based on blockchains?

References

1. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, L119:1–88, April 2016.
2. Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, et al. Hyperledger fabric: A distributed operating system for permissioned blockchains. *arXiv preprint arXiv:1801.10228*, 2018.
3. Luigi Atzori, Antonio Iera, and Giacomo Morabito. The Internet of Things: A survey. *Computer Networks*, 54:2787–2805, 2010.
4. Arshdeep Bahga and Vijay Madisetti. *Blockchain Applications: A Hands-On Approach*. VPT, 2017.
5. Thomas Baier, Claudio Di Ciccio, Jan Mendling, and Mathias Weske. Matching events and activities by integrating behavioral aspects and label analysis. *Software and System Modeling*, 17(2):573–598, 2018.
6. Thomas Baier, Jan Mendling, and Mathias Weske. Bridging abstraction layers in process mining. *Information Systems*, 46:123–139, 2014.
7. Juan Benet. Ipfs-content addressed, versioned, p2p file system, 2014.
8. Bitfury. Proof of stake vs proof of work.
9. Ethereum Blog. Merkling in Ethereum. https://blog.ethereum.org/2015/11/15/merkling-in-ethereum/. Accessed: 2018-08-17.
10. Eric A. Brewer. A certain freedom: thoughts on the CAP theorem. In Andréa W. Richa and Rachid Guerraoui, editors, *PODC*, page 335. ACM, 2010.
11. Vitalik Buterin. Ethereum: A next-generation smart contract and decentralized application platform. *URL https://github. com/ethereum/wiki/wiki/% 5BEnglish% 5D-White-Paper*, 2014.
12. Christian Cachin and Marko Vukolić. Blockchains consensus protocols in the wild. *arXiv preprint arXiv:1707.01873*, 2017.
13. Jan Camenisch and Anja Lehmann. Privacy for distributed databases via (un)linkable pseudonyms. *IACR Cryptology ePrint Archive*, 2017:22, 2017.
14. Min Chen, Shiwen Mao, and Yunhao Liu. Big Data: A Survey. *Mobile Networks and Applications*, 19(2):171–209, 2014.
15. Konstantinos Christidis and Michael Devetsikiotis. Blockchains and smart contracts for the internet of things. *IEEE Access*, 4:2292–2303, 2016.
16. Cisco. The internet of things reference model, 2014. White Paper Draft.
17. Michael Coblenz. Obsidian: A safer blockchain programming language. In *Proceedings of the 39th International Conference on Software Engineering Companion (ICSE-C)*, pages 97–99. IEEE Press, 2017.
18. E. Damaggio, R. Hull, and R. Vaculín. On the equivalence of incremental and fixpoint semantics for business artifacts with guard-stage-milestone lifecycles. *Information Systems*, 38:561–584, 2013.
19. Ali Dorri, Marco Steger, Salil S. Kanhere, and Raja Jurdak. Blockchain: A distributed solution to automotive security and privacy. *IEEE Communications Magazine*, 55(12):119–125, 2017.
20. Marlon Dumas, Marcello La Rosa, Jan Mendling, and Hajo A. Reijers. *Fundamentals of Business Process Management, Second Edition*. Springer, 2018.
21. Gilbert Fridgen, Florian Guggenmoos, Jannik Lockl, Alexander Rieger, and André Schweizer. Developing an evaluation framework for blockchain in the public sector: The ex-

ample of the german asylum process. In *Proceedings of 1st ERCIM Blockchain Workshop 2018*. European Society for Socially Embedded Technologies (EUSSET), 2018.

22 Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1994.

23 Luciano García-Bañuelos, Alexander Ponomarev, Marlon Dumas, and Ingo Weber. Optimized execution of business processes on blockchain. In *BPM'17: International Conference on Business Process Management*, Barcelona, Spain, September 2017.

24 Nitin Gaur, Luc Desrosiers, Petr Novotny, V Ramakrishna, Anthony O'Dowd, and Salman Baset. *Hands-On Blockchain with Hyperledger: Building decentralized applications with Hyperledger Fabric and Composer*. Packt Publishing Limited, 06 2018.

25 Jones Granatyr, Vanderson Botelho, Otto Robert Lessing, Edson Emílio Scalabrin, Jean-Paul Barthès, and Fabrício Enembreck. Trust and reputation models for multiagent systems. *ACM Computing Surveys (CSUR)*, 48(2):27, 2015.

26 Gideon Greenspan. Avoiding the pointless blockchain project. *Online at https://www. multichain. com/blog/2015/11/avoiding-pointless-blockchain-project*, 2015.

27 Christoph Hochreiner, Stefan Schulte, Schahram Dustdar, and Freddy Lécué. Elastic Stream Processing for Distributed Environments. *IEEE Internet Computing*, 19(6):54–59, 2015.

28 Richard Hull, Elio Damaggio, Fabiana Fournier, Manmohan Gupta, Fenno Terry Heath, Stacy Hobson, Mark Linehan, Sridhar Maradugu, Anil Nigam, Piyawadee Sukaviriya, et al. Introducing the guard-stage-milestone approach for specifying business entity lifecycles. In *International Workshop on Web Services and Formal Methods*, pages 1–24. Springer, 2010.

29 Xiaolong Jin, Benjamin W. Wah, Xueqi Cheng, and Yuanzhou Wang. Significance and challenges of big data research. *Big Data Research*, 2:59–64, 2015.

30 Audun Jøsang, Roslan Ismail, and Colin Boyd. A survey of trust and reputation systems for online service provision. *Decision support systems*, 43(2):618–644, 2007.

31 Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference*, pages 357–388. Springer, 2017.

32 S. Kiyomoto, M. S. Rahman, and A. Basu. On blockchain-based anonymized dataset distribution platform. In *2017 IEEE 15th International Conference on Software Engineering Research, Management and Applications (SERA)*, pages 85–92, June 2017.

33 Why You Shouldn't Use Airbnb: 8 Troubling Issues You Didn't Know. https://www.theinvisibletourist.com/why-you-shouldnt-use-airbnb-issues-you-didnt-know/. [Online resources].

34 Kari Korpela, Jukka Hallikas, and Tomi Dahlberg. Digital supply chain transformation toward blockchain integration. In *50th Hawaii International Conference on System Sciences*. AIS Electronic Library (AISeL), 2017.

35 Shancang Li, Li Da Xu, and Shanshan Zhao. The internet of things: a survey. *Information Systems Frontiers*, 17(2):243–259, 2015.

36 Orlenys López-Pintado, Luciano García-Bañuelos, Marlon Dumas, and Ingo Weber. Caterpillar: A blockchain-based business process management system. In *BPM'17: International Conference on Business Process Management, Demo track*, Barcelona, Spain, September 2017.

37 Orlenys López-Pintado, Luciano García-Bañuelos, Marlon Dumas, Ingo Weber, and Alexander Ponomarev. CATERPILLAR: A business process execution engine on the ethereum blockchain. Technical report, Institute of Computer Science, University of Tartu, 2018.

38 Michael Luca. Designing online marketplaces: Trust and reputation mechanisms. *Innovation Policy and the Economy*, 17(1):77–93, 2017.

39 Juri Mattila, Timo Seppälä, Catarina Naucler, Riitta Stahl, Marianne Tikkanen, Alexandra Bådenlid, Jane Seppälä, et al. Industrial blockchain platforms: An exercise in use case development in the energy industry. Technical Report 43, The Research Institute of the Finnish Economy, 2016.

40 Jan Mendling, Ingo Weber, Wil M. P. van der Aalst, et al. Blockchains for business process management - challenges and opportunities. *ACM Transactions on Management Information Systems*, 9(1):4:1–4:16, 2018.

41 Sebastien Meunier. When Do You Need Blockchain? Decision Models. https://medium.com/@sbmeunier/when-do-you-need-blockchain-decision-models-a5c40e7c9ba1, 2016.

42 Trade Ministry of Economy and Industry. Evaluation Forms for Blockchain- Based System. http://www.meti.go.jp/english/press/2017/pdf/0329_004a.pdf, 2017.

43 Miranda Mourby, Elaine Mackey, Mark Elliot, Heather Gowans, Susan E. Wallace, Jessica Bell, Hannah Smith, Stergios Aidinlis, and Jane Kaye. Are 'pseudonymised' data always personal data? implications of the gdpr for administrative data research in the uk. *Computer Law & Security Review*, 34(2):22–33, April 2018.

44 Catherine Mulligan, Jennifer Zhu Scott, Sheila Warren, and J.P. Rangaswami. Blockchain beyond the hype: A practical framework for business leaders. http://www3.weforum.org/docs/48423_Whether_Blockchain_WP.pdf, 2018.

45 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.

46 American Institute of Certified Public Accountants. *Guide to Audit Data Analytics*. AICPA, 2017.

47 Adam J. Oliner, Archana Ganapathi, and Wei Xu. Advances and challenges in log analysis. *Commun. ACM*, 55(2):55–61, February 2012.

48 There's No Middleman On This Blockchain-Based Version Of Airbnb. https://www.fastcompany.com/40524021/on-this-blockchain-based-version-of-airbnb-theres-no-middleman. [Online resources].

49 P4Titan. Slimcoin a peer-to-peer crypto-currency with proof-of-burn, 2014.

50 Morgen E Peck. Blockchain world-do you need a blockchain? this chart will tell you if the technology can solve your problem. *IEEE Spectrum*, 54(10):38–60, 2017.

51 Raghu Ramakrishnan and Johannes Gehrke. *Database Management Systems*. McGraw-Hill, Inc., New York, NY, USA, 3 edition, 2003.

52 Paul Resnick, Ko Kuwabara, Richard Zeckhauser, and Eric Friedman. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

53 Ana Reyna, Christian Martín, Jaime Chen, Enrique Soler, and Manuel Díaz. On blockchain and its integration with IoT. Challenges and opportunities. *Future Generation Computer Systems*, 88:173–190, 2018.

54 Kurt Sandkuhl and Janis Stirna. *Capability Management in Digital Enterprises*. Springer Cham, 2018.

55 Reza Shokri, Carmela Troncoso, Claudia Díaz, Julien Freudiger, and Jean-Pierre Hubaux. Unraveling an old cloak: k-anonymity for location privacy. In *WPES*, pages 115–118, 2010.

56 Yutian Sun, Wei Xu, and Jianwen Su. Declarative choreographies for artifacts. In *International Conference on Service-Oriented Computing (ICSOS)*, pages 420–434. Springer, 2012.

57 Tim Swanson. Consensus-as-a-service: a brief report on the emergence of permissioned, distributed ledger systems, 2015.

58 An Binh Tran, Qinghua Lu, and Ingo Weber. Lorikeet: A model-driven engineering tool for blockchain-based business process execution and asset management. In *BPM'18: International Conference on Business Process Management, Demo track*, Sydney, NSW, Australia, September 2018.

59 UltraNote. Absolute privacy at your fingertips.

60 Wil M. P. van der Aalst. *Process Mining - Data Science in Action, Second Edition*. Springer, 2016.

61 Wil M. P. van der Aalst, Guangming Li, and Marco Montali. Object-centric behavioral constraints. Technical report, Eindhoven University of Technology, 2017.

62 Wil M. P. van der Aalst, Mathias Weske, and Dolf Grünbauer. Case handling: a new paradigm for business process support. *Data & Knowledge Engineering*, 53(2):129–162, 2005.

63 Marko Vukolić. The quest for scalable blockchain fabric: Proof-of-work vs. bft replication. In *International Workshop on Open Problems in Network Security*, pages 112–125. Springer, 2015.

64 Marko Vukolic. Eventually returning to strong consistency. *IEEE Data Eng. Bull.*, 39(1):39–44, 2016.

65 Ingo Weber, Xiwei Xu, Régis Riveret, Guido Governatori, Alexander Ponomarev, and Jan Mendling. Untrusted business process monitoring and execution using blockchain. In *BPM'16: International Conference on Business Process Management*, Rio de Janeiro, Brazil, September 2016.

66 Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger eip-150 revision (759dccd - 2017-08-07), 2017. Accessed: 2018-01-03.

67 Karl Wüst and Arthur Gervais. Do you need a blockchain? *IACR Cryptology ePrint Archive*, 2017:375, 2017.

68 Xiwei Xu, Ingo Weber, Mark Staples, Liming Zhu, Jan Bosch, Len Bass, Cesare Pautasso, and Paul Rimba. A taxonomy of blockchain-based systems for architecture design. In *Software Architecture (ICSA), 2017 IEEE International Conference on*, pages 243–252. IEEE, 2017.

69 Ennan Zhai, David Isaac Wolinsky, Ruichuan Chen, Ewa Syta, Chao Teng, and Bryan Ford. Anonrep: Towards tracking-resistant anonymous reputation. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, pages 583–596, Santa Clara, CA, 2016. USENIX Association.

70 Kaiwen Zhang and Hans-Arno Jacobsen. Towards dependable, scalable, and pervasive distributed ledgers with blockchains. In *38th IEEE International Conference on Distributed Computing Systems*, pages 1337–1346. IEEE Computer Society, 2018.

71 Kaiwen Zhang, Roman Vitenberg, and Hans-Arno Jacobsen. Deconstructing blockchains: Concepts, systems, and insights. In *Proceedings of the 12th ACM International Conference on Distributed and Event-based Systems*, pages 187–190. ACM, 2018.

Participants

- Michael Coblenz
Carnegie Mellon University – Pittsburgh, US
- Søren Debois
IT University of Copenhagen, DK
- Claudio Di Ciccio
Wirtschaftsuniversität Wien, AT
- Alevtina Dubovitskaya
EPFL – Lausanne, CH
- Marlon Dumas
University of Tartu, EE
- Fabiana Fournier
IBM – Haifa, IL
- Avigdor Gal
Technion – Haifa, IL
- Luciano García-Bañuelos
University of Tartu, EE
- Stephan Haarmann
Hasso-Plattner-Institut – Potsdam, DE
- Richard Hull
IBM TJ Watson Research Center – Yorktown Heights, US
- Hans-Arno Jacobsen
TU München, DE
- Mieke Jans
Hasselt University, BE
- Agnes Koschmider
KIT – Karlsruher Institut für Technologie, DE
- Qinghua Lu
Data61, CSIRO – Sydney, AU
- Raimundas Matulevičius
University of Tartu, EE
- Jan Mendling
Wirtschaftsuniversität Wien, AT
- Petr Novotny
IBM TJ Watson Research Center –Yorktown Heights, US
- Sooyong Park
Sogang University – Seoul, KR
- Stefanie Rinderle-Ma
Universität Wien, AT
- Stefan Schulte
TU Wien, AT
- Jerome Simeon
Clause Inc. – New York, US
- Ludwig Stage
Tübingen, DE
- Mark Staples
Data61, CSIRO – Eveleigh, AU
- Barbara Weber
Technical University of Denmark – Lyngby, DK
- Ingo Weber
Data61, CSIRO – Sydney, AU
- Francesca Zerbato
University of Verona, IT
- Kaiwen Zhang
ETS – Montreal, CA

Report from Dagstuhl Seminar 18341

Formalization of Mathematics in Type Theory

Edited by

Andrej Bauer[1], Martín H. Escardó[2], Peter L. Lumsdaine[3], and Assia Mahboubi[4]

1 **University of Ljubljana, SI,** andrej.bauer@fmf.uni-lj.si
2 **University of Birmingham, GB,** m.escardo@cs.bham.ac.uk
3 **University of Stockholm, SE,** p.l.lumsdaine@math.su.se
4 **INRIA – Nantes, FR,** assia.mahboubi@inria.fr

Abstract

Formalized mathematics is mathematical knowledge (definitions, theorems, and proofs) represented in digital form suitable for computer processing. The central goal of this seminar was to identify the theoretical advances and practical improvements needed in the area of formalized mathematics, in order to make it a mature technology, truly useful to a larger community of students and researchers in mathematics. During the seminar, various software systems for formalization were compared, and potential improvements to existing systems were investigated. There have also been discussions on the representation of algebraic structures in formalization systems.

Seminar August 19–24, 2018 – http://www.dagstuhl.de/18341
2012 ACM Subject Classification Theory of computation → Logic and verification, Theory of computation → Type theory
Keywords and phrases formal methods, formalized mathematics, proof assistant, type theory
Digital Object Identifier 10.4230/DagRep.8.8.130
Edited in cooperation with Auke B. Booij

1 Executive Summary

Andrej Bauer
Martín H. Escardó
Peter L. Lumsdaine
Assia Mahboubi

We and all the participants were delighted to benefit from Dagstuhl's inspiring environment.

Proof assistants are receiving increased attention from users with a background in mathematics, as opposed to their traditional users from theoretical computer science/logic/program verification, and this was the major focus of the meeting. This is true in particular of proof assistants based on dependent types, probably due in part to the advent of homotopy type theory, developed in the proof assistants Coq, Agda and Lean.

The audience of the seminar was thus rather unusual in composition, and featured several experienced researchers used to attending seminars at the Mathematisches Forschungsinstitut Oberwolfach, and visiting Schloss Dagstuhl for the first time. In order to foster discussion and fuse collaborations, we adopted a different format from the standard string of slide-based talks: talks in the morning, so that people get to know the work of each other, and working

in groups in the afternoon. At the end of each day, before dinner, each group presented a summary of the outcomes of their meetings to all participants, which allowed inter-group discussion and collaboration. This had been tried before by some of the organizers, in the course of Dagstuhl seminar 16112, and worked just as well in our case.

Working group topics were proposed by the audience on the first day, by giving short presentations of a few minutes and writing topics in the board. Some were quite specialized and homogeneous (e.g. the cubical type theory group), and allowed people to have a focussed collaborative brainstorming on a specific open problem of the field. Some were more open-ended, and allowed people to confront various approaches to the same issue/concept in different systems (different proof assistants, computer algebra systems, etc.).

Some people did applied work, such as trying to compute the so-called Brunerie number from an existing proof in homotopy type theory, in order to identity and fix inefficiency problems in proofs assistants based on cubical type theory. Some people used their spare time to solve the "Dagstuhl dinner" problem. Details of the topics discussed are in the reports produced by each group.

This was a rather productive meeting, and people from different scientific backgrounds not only met but talked together effectively, solving and identifying problems to work on collaboratively in future.

2 Table of Contents

3 Overview of Talks

3.1 Deriving on Steroids – for proof assistants

Jacques Carette (McMaster University – Hamilton, CA)

Joint work of Jacques Carette, Russell O'Connor
Main reference Jacques Carette, Russell O'Connor: "Theory Presentation Combinators", in Proc. of the Intelligent Computer Mathematics – 11th International Conference, AISC 2012, 19th Symposium, Calculemus 2012, 5th International Workshop, DML 2012, 11th International Conference, MKM 2012, Systems and Projects, Held as Part of CICM 2012, Bremen, Germany, July 8-13, 2012. Proceedings, Lecture Notes in Computer Science, Vol. 7362, pp. 202–215, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-31374-5_14

Developing large theory graphs is a lot of work – even without the proofs. It turns out that the contents of mathematics is highly structured, and that structure can be used to alleviate the development of theories. This naturally leads to theory presentation combinators, which transform existing knowledge into new, in a scalable manner.

Experience shows that 3 main combinators arise naturally: extension, renaming, and combination. Extending adds a new concept to a theory; here we use the *tiny theories* approach, which is to always add a single concept at a time. Renaming is necessary as mathematical conventions for things which are "the same" nevertheless use different symbols in different contexts. First-class renaming allows this "sameness" to be tracked automatically. Lastly, combination is a generalization of union which takes care of necessary gluings when we want to merge two theories with a common ancestry. In other words, for theory presentations, the "diamond problem" is actually a blessing.

This approach appears to scale well. It also has a denotational theory that is quite familiar, as it re-uses the category of contexts, fibrations and pullbacks as its main ingredients.

From there, the ideas of "deriving" from Haskell really kick in: it is straightforward to notice that many constructions from Universal Algebra lift immediately to this setting. Thus one can automatically derive new theories, such as that of homomorphisms, from existing theories. Further more, one can also derive term languages and accompanying functions, automatically from the signature of a theory. This can also be staged, so that meta-programming comes into scope, so that simplistic optimizing compilers for terms of a theory's language can be automatically derived.

3.2 Classical Analysis with Coq

Cyril Cohen (INRIA Sophia Antipolis, FR) and Assia Mahboubi (INRIA – Nantes, FR)

Joint work of Reynald Affeldt, Cyrill Cohen, Damien Rouhling, Assia Mahboubi, Pierre-Yves Strub

In this talk I presented an ongoing effort to develop a Coq formal library, MathComp-Analysis [1], about classical real analysis. Almost all existing proof assistants on the market have been used to investigate the formalization of real, and sometimes also complex, analysis. A survey by Boldo et al. reviews the different approaches and the breadth of the existing developments [2].

Our motivation for designing yet another formal analysis library is twofold. First, we rely on strong classical axioms, so as to get closer to the logical formalism used in classical

mathematics. Second, we design it along the formalization methodology put into practice in the MATHEMATICAL COMPONENTS libraries [3]. The latter libraries are essentially geared towards algebra and this work aims at providing an extension for topics in analysis.

The main original contributions lie in the effort put in the infrastructure of MATHCOMP-ANALYSIS: automation, notations, etc... I presented more in details two mechanisms to do asymptotic reasoning: one to simplify proofs about filters and another to deal with Bachmann-Landau notations.

References

1 Reynald Affeldt, Cyril Cohen, Assia Mahboubi, Damien Rouhling, and Pierre-Yves Strub. Analysis library compatible with Mathematical Components. https://github.com/math-comp/analysis/releases/tag/0.1.0 (last accessed: 2018/10/01), 2018.
2 Sylvie Boldo, Catherine Lelay, and Guillaume Melquiond. Formalization of real analysis: a survey of proof assistants and libraries. *Mathematical Structures in Computer Science*, 26(7):1196–1233, 2016.
3 Assia Mahboubi and Enrico Tassi. *Mathematical Components*. Available at: https://math-comp.github.io/mcb/, 2016. With contributions by Yves Bertot and Georges Gonthier.

3.3 Isabelle/HOL Demo

Manuel Eberl (TU München, DE)

I demonstrated the interactive theorem prover Isabelle in the logic HOL by showing how to prove the infinitude of primes in it. I also showed some of its more specialized tactics, like those for approximation of real numbers or real limits, and the code generation feature.

3.4 A Coq Formalization of Digital Filters

Diane Gallois-Wong (Laboratoire de Recherche en Informatique – Orsay, FR)

Joint work of Boldo, Sylvie; Hilaire, Thibault
Main reference Diane Gallois-Wong, Sylvie Boldo, Thibault Hilaire: "A Coq Formalization of Digital Filters", in Proc. of the Intelligent Computer Mathematics – 11th International Conference, CICM 2018, Hagenberg, Austria, August 13-17, 2018, Proceedings, Lecture Notes in Computer Science, Vol. 11006, pp. 87–103, Springer, 2018.
URL http://dx.doi.org/10.1007/978-3-319-96812-4_8

Digital filters are small iterative algorithms, used as basic bricks in signal processing and control systems. Therefore, they have numerous application domains, including communication, automotive, robotics, aeronautics, etc. They are usually studied as mathematical objects using real numbers. However, to be used in practice, they need to be implemented, which implies finite precision arithmetic (floating- or fixed-point numbers) and rounding errors. Moreover, propagation of these rounding errors through iteration makes these errors potentially critical but also hard to study. That is why we aim at providing a formal analysis of the rounding errors in digital filters, using the Coq proof assistant. In our current formalization, we define three algorithms used to implement digital filters, called realizations.

We prove that they are equivalent, so that we can focus on one of them for the rest of the error analysis. Then, we formally prove two theorems that are essential to the error analysis: the theorem of the error filter, that characterizes the final error between the implemented filter and the ideal one using infinite precision, and the Worst-Case Peak-Gain theorem, that bounds the output corresponding to a bounded input.

3.5 An overview of UniMath

Daniel R. Grayson (Urbana, US)

In this 30 minute talk we give a brief overview of UniMath, the formalization project started by Voevodsky, Ahrens, and me, that aims to formalize a substantial body of mathematics in the univalent foundations, building on the original formalization by Voevodsky from 2009, the "Foundations". The code (174K lines) is in Coq and is hosted at http://unimath.org/. My personal goal for the next year is to formalize a preprint of mine on algebraic K-theory, so when I finally submit it for publication, I can include the formalization to make the job of the referee easier.

We give a tour of the code, showing how interaction with it works in ProofGeneral in emacs, touching upon univalence, the implementation of groups, and Voevodsky's resizing axioms.

3.6 Formalization of Smooth Manifolds in Isabelle/HOL

Fabian Immler (Carnegie Mellon University – Pittsburgh, US) and Bohua Zhan (Chinese Academy of Sciences – Beijing, CN)

We formalize the definition and basic properties of smooth manifolds in Isabelle/HOL. Concepts covered include partition of unity, tangent and cotangent spaces, and the fundamental theorem for line integrals. We also construct some concrete manifolds such as spheres and projective spaces. The formalization makes extensive use of the existing libraries for topology and analysis. The existing library for linear algebra is not flexible enough for our needs. We therefore set up the first systematic and large scale application of "types to sets". It allows us to automatically transform the existing (type based) library of linear algebra to one with explicit carrier sets.

3.7 Heuristics for rewrite search

Scott Morrison (Australian National University – Canberra, AU)

Joint work of Scott Morrison, Keeley Hoek

I gave a demo of my recent work formalizing category theory in Lean, both discussing my goals as mathematician visiting the world of interactive theorem proving, and showing off some fun graph visualizations of an algorithm for proving equational lemmas using rewriting guiding by edit distance heuristics.

In general, I have been trying to understand how far writing mathematics in a modern interactive theorem prover is from the usual experience of writing and explaining mathematics to other humans. (Of course, the answer for now is "too far".) Category theory is an interesting and easy test case, as frequently in proofs and constructions there are quite considerable verifications which ought to be undertaken (checking functors are functorial, natural transformations natural) but which are very frequently omitted in human mathematics. I've been trying to write a category theory library working within the constraint that none of these verifications may be performed with human assistance (and ideally, should be kept entirely out of human sight). Of course, this requires developing at the same time a certain amount of automation particular to the domain of the mathematics being formalized. It is essentially for this reason that I've chosen to work in Lean: it seems to have the most flexible and easy to learn mechanism for writing new automation amongst modern theorem provers.

Do we hope one day to have a non-trivial portion of research mathematics performed with the aid of computers? (Here I mean the actual research, not merely post hoc formalization.) If so, I think it will be necessary that 'writing tactics' becomes easy enough that it is within reach of end users, not just developers of the interactive theorem provers. For now, of course, it is not easy enough, but I have been encouraged by working in Lean, and observing my (mathematics) students coming to grips with Lean and writing tactics in Lean. (The biggest obstacle may just be that nearly all mathematicians, and still most mathematics students, aren't at all familiar with functional programming and working with monads! Dependent type theories themselves are no obstacle.)

In my demo I showed two related recent pieces of work. One was an algorithm for proving equational goals via rewriting, using heuristics based on edit distance to explore the graph of possible rewrites by a given set of lemmas. (I think the audience enjoyed the graphical visualizations of the proof searches!) Along with a student Keeley Hoek, we're now incorporating classification techniques, using a support vector machine to dynamically reweight the tokens appearing in expressions as the search proceeds. This is early work, but tentatively it appears that this can help focus the search on the key steps, avoiding needlessly exploring minor irrelevant rewrites. We're new to this field, and hoping to learn more about previous work in this direction, and especially hoping to come up with heuristics for generating "stepping stone" intermediate goals based on analyzing partial search graphs.

The second was an illustration of how this algorithm can elide many of the "boring" proofs in a basic category theory library. I quickly showed some examples of proofs of the Yoneda lemma from other interactive theorem provers, followed by the very short proofs in my library in Lean. These successfully rely on some basic automation, and the heuristics for rewrite searches described above, to allow us to just write the statements a mathematician

would write, omitting all the easy verifications. As an example, we can reduce the entire definition of the Yoneda functor itself to

```
def yoneda : C ⇒ ((C^op) ⇒ (Type v₁)) := λ' X, λ' Y : C, Y → X
```

with two functoriality and one naturality statement being synthesized behind the scenes. (I found some formalizations of this statement that occupied more than a page.) Obviously this is an extreme example, but it illustrates my goal that automation should strive to meet the mathematician, rather than the other way round, when possible.

Participating in the Dagstuhl seminar was really exciting for me – it was a great opportunity to make contact with the community around interactive theorem provers, and it was great that mathematicians new to the field were made so welcome!

3.8 Cubical Agda Demo

Anders Mörtberg (Chalmers University of Technology – Göteborg, SE)

In this short demo I showed a cubical version of the Agda proof assistant implemented by Andrea Vezzosi. This system allows for a direct proof of functional extensionality and also the univalence axiom. Another exciting aspect is that it allows the user to define higher inductive types with good computational behavior. This was illustrated by live proving that the torus is equivalent to the product of two circles.

3.9 Semi-formal verification as a routine tool

Neil Strickland (University of Sheffield, GB)

Proof assistants are rarely used by working mathematicians for new research. Many people have thought about what proof assistants can currently do, and how one should work from there towards what mathematicians need. In this talk we look at the problem from a different angle. There are other mathematical software systems that are very widely used by researchers, including Sage, Mathematica and Maple. In particular, the author has made extensive use of Maple for a kind of semi-formal verification of some kinds of mathematical arguments. In this talk we describe this experience, and discuss how to narrow the gap with fully formal verification. It would be helpful if proof assistants could interface in some way with systems such as Maple, and we also discuss some issues related to this.

3.10 Formal Abstracts

Floris van Doorn (University of Pittsburgh, US)

Joint work of Tom Hales, Floris van Door
URL https://github.com/formalabstracts/formalabstracts

Formal Abstracts is a ambitious project to build a large database of theorems and definitions from almost all branches of mathematics in both human-readable and machine-readable form. It will serve as a database for machine-learning projects, can be used to semantically search the database of mathematical definitions and theorems, could be used for exploring mathematics, and to translate mathematics between languages, preserving the semantics. In this talk I gave an overview of the goals and concrete plans of the Formal Abstracts project.

3.11 Structuring principles for specifications in Isabelle

Makarius Wenzel (Augsburg, DE)

URL http://files.sketis.net/Dagstuhl2018/Isabelle_Structure.tar.gz

This is a brief overview of Local Theory Specifications in Isabelle/Pure, which are extensively used in Isabelle/HOL libraries and applications: unnamed contexts, locales, type classes. The included theory document (Isabelle_Structure.tar.gz) is for Isabelle2018.

3.12 The Isabelle Prover IDE after 10 years of development

Makarius Wenzel (Augsburg, DE)

URL https://sketis.net/wp-content/uploads/2018/08/Dagstuhl2018.pdf

The main ideas around Isabelle/PIDE go back to summer 2008. This is an overview of what has been achieved in the past 10 years, with some prospects for the future. Where can we go from here as Isabelle community? (E.g. towards alternative front-ends like Visual Studio Code; remote prover sessions "in the cloud"; support for collaborative editing of large formal libraries.) Where can we go as greater ITP community (Lean, Coq, HOL family)?

4 Working groups

4.1 Dagstuhl's Happy Diner Problem

Auke Booij (University of Birmingham, GB) and Floris van Doorn (University of Pittsburgh, US)

URL https://github.com/fpvandoorn/Dagstuhl-tables/

We have investigated Dagstuhl's Happy Diner Problem to find optimal seating arrangements during the meals at Dagstuhl.

The problem statement: What is the minimum number of meals so that each of the n conference participants can share at least one meal with every other participant when eating at tables of at most k persons? We call this number $T(n, k)$.

In particular, we have an unlimited number of tables, and we do not require that any two participants have a meal together exactly once, or that every table is fully occupied.

During the seminar, we have made progress on this problem using various techniques. This work is being documented via Github [1], and is ongoing.

- We have found several relations between various entries in the table of values $T(n, k)$, yielding both lower bounds and upper bounds for many entries. These relations allow us to fill in many entries in the table without any further exhaustive searches.
- We have manually computed certain entries $T(n, k)$, allowing us to fill in certain regions of the table.
- We collaborated with Michael Trott to use Mathematica's built-in SAT solver to find upper bounds for $T(n, k)$ for certain values of n and k.
- We have compared this problem with various related problems, such as the Oberwolfach problem [2], the Social Golfer problem [3, 4], and finding Kirkman Triple Systems[5]. In some cases, this allowed us to find values $T(n, k)$.

As a result of the work, we have submitted sequences to the Online Encyclopedia of Integer Sequences: A318240 and A318241. We have summarized the results in Table 1.

References

1 Floris P. van Doorn, Auke B. Booij. *Dagstuhl's Happy Diner problem*. Available at: https://github.com/fpvandoorn/Dagstuhl-tables/.
2 Sarah Holliday. *Sarah's Oberwolfach Problem Page*. Available at: http://facultyweb.kennesaw.edu/shollid4/oberwolfach.php. Accessed October 2018.
3 Social Golfer Problem on Wolfram MathWorld. Available at: http://mathworld.wolfram.com/SocialGolferProblem.html. Accessed October 2018.
4 A107431 on the Online Encyclopedia of Integer Sequences. Available at: https://oeis.org/A107431. Accessed October 2018.
5 Dijen K. Ray-Chaudhuri, Richard M. Wilson. *Solution of Kirkman's schoolgirl problem*. In Proc. of Symp. in Pure Math, Vol 19, 1971.

Table 1 Table of solutions $T(n,k)$, or ranges of possible solutions. Bold numbers are optimal solutions in the sense that every conference participants shares a meal with every other participant *exactly* once.

n / k	2	3	4	5	6	7	8
1	**0**	**0**	**0**	**0**	**0**	**0**	**0**
2	**1**	1	1	1	1	1	1
3	3	**1**	1	1	1	1	1
4	**3**	3	**1**	1	1	1	1
5	5	3	3	**1**	1	1	1
6	**5**	4	3	3	**1**	1	1
7	7	4	3	3	3	**1**	1
8	**7**	4	3	3	3	3	**1**
9	9	**4**	4	3	3	3	3
10	**9**	6	4	4	3	3	3
11	11	6	5	4	3	3	3
12	**11**	6	5	4	3	3	3
13	13	7	5	5	4	3	3
14	**13**	7	5	5	4	4	3
15	15	**7**	5	5	4	4	3
16	**15**	9	**5**	5	4	4	3
17	17	9	6-9	5	4	4	3-4
18	**17**	9	7-9	5-6	4	4	3-4
19	19	10	7-9	5-6	5-6	4	3-4
20	**19**	10	7-9	5-6	5-6	4-6	4
21	21	**10**	8-9	6	5-6	4-6	4-5
22	**21**	12	8-9	6	5-6	4-6	4-5
23	23	12	8-9	6	5-6	4-6	4-5
24	**23**	12	8-9	6	5-6	5-6	4-5
25	25	13	9	**6**	6	5-6	4-5
26	**25**	13	9	7-9	6	5-6	4-5
27	27	**13**	9	7-9	6-7	5-6	5
28	**27**	15-16	**9**	8-9	6-7	5-7	5
29	29	15-16	10-11	8-9	6-7	5-7	5
30	**29**	15-16	11	8-11	6-7	5-7	5

4.2 Interoperability of systems

Mario Carneiro (Carnegie Mellon University – Pittsburgh, US), Gaëtan Gilbert (INRIA – Nantes, FR), Fabian Immler (Carnegie Mellon University – Pittsburgh, US), Maria Emilia Maietti (University of Padova, IT), and Makarius Wenzel (Augsburg, DE)

There was a total of 4 sessions, with slightly varying participants. Some notable topics of discussion:

- General problems of adjusting the logical languages, e.g. Lean vs. HOL, or other Type Theories.
- Questions about alignment of library content, e.g. the translated version of Nat vs. the existing one in the target system.

- An unorthodox approach to import HOL4 + CakeML into Isabelle, based on Isabelle/ML "virtualization" and replay of the original HOL4 theory and proof scripts directly in Isabelle/HOL (backed by concrete experiments by Fabian Immler).

Immler and Wenzel later continued the prototype of "virtual HOL4 inside Isabelle"; this work is likely to become part of future releases of any of these proof assistants.

4.3 Debugging Coq

Gaëtan Gilbert (INRIA – Nantes, FR) and Daniel R. Grayson (Urbana, US)

Gaëtan Gilbert and I looked into the internals of Coq to try to figure out how to make the resizing axioms work in UniMath without disabling all universe checking using the "type in type" option. The main result was a succinct bug report to the Coq team at INRIA, which we hope will be acted upon soon. I also learned from Gaëtan some tricks for debugging Coq in the OCaml debugger, which will help if I have to look more deeply into the problem.

4.4 Structures

Assia Mahboubi (INRIA – Nantes, FR), Yves Bertot (INRIA Sophia Antipolis, FR), Jacques Carette (McMaster University – Hamilton, CA), Cyril Cohen (INRIA Sophia Antipolis, FR), Diane Gallois-Wong (Laboratoire de Recherche en Informatique – Orsay, FR), Georges Gonthier (INRIA Saclay – Île-de-France, FR), Florent Hivert (Laboratoire de Recherche en Informatique – Orsay, FR), Johannes Hölzl (Free University Amsterdam, NL), Scott Morrison (Australian National University – Canberra, AU), Russell O'Connor (Blockstream – Montreal, CA), Claudio Sacerdoti Coen (University of Bologna, IT), Bas Spitters (Aarhus University, DK), Michael Trott (Wolfram Research – Champaign, US), Hoang Le Truong (Universität des Saarlandes, DE), Josef Urban (Czech Technical University – Prague, CZ), and Makarius Wenzel (Augsburg, DE)

The participants of this working group have discussed the representation of algebraic structures both in computer algebra systems and in a collection of different proof assistants. Several short presentations have fostered the discussions, including:

- Canonical Structures in MathComp by Georges Gonthier;
- Structures in Sage by Florent Hivert (based on excerpts of Nicolas M. Thiéry's talks at CICM, July 28th of 2016, Bialystok);
- MathClasses by Bas Spitters;
- Unification hints by Claudio Sacerdoti;
- Type classes/Locales in Isabelle by Makarius Wenzel;
- Type classes in Lean by Johannes Hölzl;
- Soft typing in Mizar by Josef Urban;
- auto2, by Bohua Zhan's.

The objective was to review the different solutions, their assets and their limitations. In fact, a significant part of the discussion was devoted to identifying the requirements of the various stakeholders-roles, namely system builders, library builders, advanced users and users. Making more precise these requirements should help benchmarking the different ways of designing and implementing a graph of structures in a proof assistant, together with the related tools for inference, search, debug, etc.

The second main outcome of the working group is a list of known difficult problems related to the inference of instances of algebraic structures, in various contexts. The items in this list are very diverse in nature, ranging from the algorithmic issues in the inference algorithms implemented by proof assistants, to the design of complex hierarchies like ordered algebraic structures, and to the cost of changing the representation of objects (e.g. dense vs sparse polynomials).

4.5 Cubical Working Group

Anders Mörtberg (Chalmers University of Technology – Göteborg, SE), Carlo Angiuli (Carnegie Mellon University – Pittsburgh, US), Guillaume Brunerie (University of Stockholm, SE), Kuen-Bang (Favonia) Hou (University of Minnesota – Minneapolis, US), Simon Huber (University of Göteborg, SE), Dan Licata (Wesleyan University – Middletown, US), Ian Orton (University of Cambridge, GB), Bas Spitters (Aarhus University, DK), and Jonathan Sterling (Carnegie Mellon University – Pittsburgh, US)

The members of the cubical working group worked on a variety of problems related to cubical type theories. These theories provide computational justifications to the univalence axiom and higher inductive types and there are now multiple implementations based on these new type theories. Many of the authors of these systems were present at the meeting which led to very fruitful collaborations among experts that are not often at the same place – thanks to the organizers and Dagstuhl for providing us with this opportunity.

The main problem that we worked on was to better understand the various computational inefficiencies that seem present in all of the implementations of cubical type theories. For example we noticed that the computation time and memory usage was heavily dependent on how loops were nested when computing winding numbers. With these examples we could benchmark the various systems and get new ideas for how to optimize the particular systems. This led to a variety of new optimizations which increased the performance on multiple of the examples.

We also optimized the proof of one of the key lemmas underlying the most complicated part of the algorithms in all of these systems. In this algorithm we only need a special case of a general lemma and we found a new proof of this special case, which we now call the "Dagstuhl lemma". The lemma was used to optimize the implementation of the cubicaltt proof checker and it was also formally verified in Agda during the meeting.

One of the major open problems in implementing cubical type theory is to define a simple and efficient notion of evaluation which is adequate for open terms. The presence of the diagonal cofibrations in cartesian cubical type theory complicates the question and constrains the potential solutions, leading to a form of evaluation which is executed relative to an evolving equational theory on dimensions. During the seminar, members of the cubical type

theory working group collaborated to work out the invariants and operations of a semantic domain for open computation, which will form the backbone of the algorithm to decide definitional equivalence in implementations of cartesian cubical type theory, such as the redtt proof assistant.

4.6 Subjecting mathematicians to proof assistants

Neil Strickland (University of Sheffield, GB), Sophie Bernard (INRIA Sophia Antipolis, FR), Auke Booij (University of Birmingham, GB), Mario Carneiro (Carnegie Mellon University – Pittsburgh, US), Manuel Eberl (TU München, DE), Martín H. Escardó (University of Birmingham, GB), Daniel R. Grayson (Urbana, US), Nicolai Kraus (University of Nottingham, GB), Scott Morrison (Australian National University – Canberra, AU), Anja Petkovic (University of Ljubljana, SI), Makarius Wenzel (Augsburg, DE), and Bohua Zhan (Chinese Academy of Sciences – Beijing, CN)

URL http://neil-strickland.staff.shef.ac.uk/dagstuhl/

The aim of this working group was to develop examples and documentation explaining the use of proof assistants to working mathematicians. The emphasis is on issues likely to arise when trying to formalize new research, and issues where proof assistants fit poorly with a mathematician's natural expectations and intuitions. During the meeting we gathered a lot of useful information, from presentations of code as well as discussions of conceptual issues. Since the meeting a substantial amount of work has been done towards assembling this information into a useful set of web pages, and the Lean community has also contributed further code and advice. This work is still ongoing.

Our discussions in the working group were organized around the four tasks described below. Some solutions were presented by Sophie Bernard (Coq + ssreflect), Manuel Eberl (Isabelle-HOL) and Bohua Zhan (Isabelle-FOL), and the detailed walk-through of this code was very illuminating. There were also conversations outside the working group in which various people explained useful things; thanks are especially due to Mario Carneiro, Scott Morrison and Makarius Wenzel.

In brief, the tasks were as follows.

- Prove that for any natural number n, there is a prime p with $p > n$.
- Set up the theory of the group of units in a commutative ring.
- Set up the theory of the ideal of nilpotent in a commutative ring, and the corresponding quotient ring.
- Set up the theory of chained preorders (a kind of discrete combinatorial structure on a finite set).

The detailed specification of the tasks covers a number of issues that may cause difficulty: locating standard results in the standard library, confusion between different implementations of the natural numbers, the general framework for abstract algebra, subobjects and quotient objects, finiteness and decidability, and so on.

Participants

- Benedikt Ahrens
University of Birmingham, GB
- Carlo Angiuli
Carnegie Mellon University – Pittsburgh, US
- Andrej Bauer
University of Ljubljana, SI
- Sophie Bernard
INRIA Sophia Antipolis, FR
- Yves Bertot
INRIA Sophia Antipolis, FR
- Auke Booij
University of Birmingham, GB
- Guillaume Brunerie
University of Stockholm, SE
- Jacques Carette
McMaster University – Hamilton, CA
- Mario Carneiro
Carnegie Mellon University – Pittsburgh, US
- Cyril Cohen
INRIA Sophia Antipolis, FR
- Manuel Eberl
TU München, DE
- Martín H. Escardó
University of Birmingham, GB
- Diane Gallois-Wong
Laboratoire de Recherche en Informatique – Orsay, FR
- Gaëtan Gilbert
INRIA – Nantes, FR
- Georges Gonthier
INRIA Saclay – Île-de-France, FR
- Daniel R. Grayson
Urbana, US
- Philipp Haselwarter
University of Ljubljana, SI
- Florent Hivert
Laboratoire de Recherche en Informatique – Orsay, FR
- Johannes Hölzl
Free University Amsterdam, NL
- Kuen-Bang (Favonia) Hou
University of Minnesota – Minneapolis, US
- Simon Huber
University of Göteborg, SE
- Fabian Immler
Carnegie Mellon University – Pittsburgh, US
- Nicolai Kraus
University of Nottingham, GB
- Dan Licata
Wesleyan University – Middletown, US
- Peter L. Lumsdaine
University of Stockholm, SE
- Assia Mahboubi
INRIA – Nantes, FR
- Maria Emilia Maietti
University of Padova, IT
- Anders Mörtberg
Chalmers University of Technology – Göteborg, SE
- Scott Morrison
Australian National University – Canberra, AU
- Russell O'Connor
Blockstream – Montreal, CA
- Ian Orton
University of Cambridge, GB
- Anja Petkovic
University of Ljubljana, SI
- Claudio Sacerdoti Coen
University of Bologna, IT
- Bas Spitters
Aarhus University, DK
- Jonathan Sterling
Carnegie Mellon University – Pittsburgh, US
- Neil Strickland
University of Sheffield, GB
- Michael Trott
Wolfram Research – Champaign, US
- Hoang Le Truong
Universität des Saarlandes, DE
- Josef Urban
Czech Technical University – Prague, CZ
- Floris van Doorn
University of Pittsburgh, US
- Makarius Wenzel
Augsburg, DE
- Bohua Zhan
Chinese Academy of Sciences – Beijing, CN

Report from Dagstuhl Seminar 18351

Modeling for Sustainability

Edited by

Gordon Blair[1], Betty H. C. Cheng[2], Lorenz Hilty[3], and Richard F. Paige[4]

1 **Lancaster University, GB,** gordon.s.blair@googlemail.com
2 **Michigan State University – East Lansing, US,** chengb@cse.msu.edu
3 **Universität Zürich, CH,** hilty@ifi.uzh.ch
4 **University of York, GB,** richard.paige@york.ac.uk

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 18351 "Modeling for Sustainability".

Seminar August 26–31, 2018 – http://www.dagstuhl.de/18351
2012 ACM Subject Classification Social and professional topics → Sustainability, Computing methodologies → Model development and analysis, Computing methodologies → Modeling methodologies
Keywords and phrases modeling for sustainability, sustainability dimensions, environmental sustainability, social sustanability, economic sustainability, model driven engineering
Digital Object Identifier 10.4230/DagRep.8.8.146
Edited in cooperation with Sedef Akinli Kocak

1 Executive Summary

Gordon Blair
Betty H. C. Cheng
Lorenz Hilty
Richard F. Paige

Many different kinds of models, from engineering models to scientific models, have to be integrated and coordinated to support sustainability systems such as smart grid or cities, i.e., dynamically adaptable resource management systems that aim to improve the techno-economic, social, and environmental dimensions of sustainability. Scientific models help understand sustainability concerns and evaluate alternatives, while engineering models support the development of sustainability systems. As the complexity of these systems increases, many challenges are posed to the computing disciplines to make data and model-based analysis results more accessible as well as integrate scientific and engineering models while balancing trade-offs among varied stakeholders. This seminar explored the intrinsic nature of both scientific and engineering models, the underlying differences in their respective foundations, and the challenges related to their integration, evolution, analysis, and simulation including the exploration of what-if scenarios.

Sustainability systems must provide facilities for the curation and monitoring of data sets and models and enable flexible (open) data and model integration, e.g., physical laws, scientific models, regulations and preferences, possibly coming from different technological

foundations, abstractions, scale, technological spaces, and world views. This also includes the continuous, automated acquisition and analysis of new data sets, as well as automated export of data sets, scenarios, and decisions. The main function is to support the generation of what-if scenarios to project the effects on the different sustainability dimensions, and support the evaluation of externalities, especially for non rapidly renewable resources. Since the predictions are necessarily probabilistic, the system must be able to assess the uncertainty inherent in all its actions and provide suitable representations of uncertainty understandable by users. In addition to generating what-if scenarios to explore alternate model instantiations, the tool should be capable of generating suggestions for how to reach user-specified goals including quantifiable impacts and driving the dynamic adaptation of sustainability systems. These powerful services must be made accessible to the population at large, regardless of their individual situation, social status, and level of education.

This seminar explored how Model-Driven Engineering (MDE) will help to develop such an approach, and in particular i) how modeling frameworks would support the integration of the various heterogeneous models, including both engineering and scientific models; ii) how domain specific languages (DSLs) would (a) support the required socio-technical coordination, i.e., engage engineers, scientists, decision makers, communities, and the general public; and (b) integrate analysis/probabilistic/user models into the control loop of smart CPS (cyber physical system). DSLs are also supposed to provide the right interface (in terms of abstractions/constructs) to be used as tools for discovering problems and evaluating ideas.

The seminar served to identify critical disciplines and stakeholders to address MDE for sustainability and the research roadmap of the MDE community with regards to the development of sustainability systems. In particular, the seminar identified and explored four key areas: 1) research challenges relevant to modeling for sustainability (M4S); 2) a multidisciplinary collection of relevant literature to provide the foundation for exploring the research challenges; 3) three case studies from different application domains that provide a vehicle for illustrating the M4S challenges and for validating relevant research techniques; and 4) the human and social aspects of M4S.

The cumulative results of the work performed at the seminar and subsequent collaborations will help to establish the required foundations for integrating engineering and scientific models, and to explore the required management facilities for evaluating what-if scenarios and driving adaptive systems. In addition, we envision to produce as an outcome of the seminar a representative case study that will be used by the community to assess and validate contributions in the field of modeling for sustainability.

2 Table of Contents

3 Overview of Talks

3.1 Models of and for Sustainability in my domain

Lucy Bastin (Aston University – Birmingham, GB)

My background is multi-disciplinary – from zoologist to GIS software developer to researcher at the policy interface, which has covered pretty much all the types of models that have been mentioned so far. I will focus on some additional sustainability definitions from my current work which may be relevant to our discussions. The first is classic environmental "sustainability" whose ultimate goal is SDGs, aimed at sustaining human life on earth at a certain quality. We aim to conserve biodiversity and ecosystem services by sharing benefits / reducing human/wildlife conflict. This involves

- observing, sampling and modelling that biodiversity,
- inference / transformation to estimate the ecosystem services it supports,
- prediction of the likely human actions, landscape modifications and movements that will affect biodiversity,
- prediction of the ways that people and wildlife will respond to climate and infrastructure changes,
- multi-objective planning with diverse stakeholders,
- identification of the pinch points in the landscape where support will be needed.

The second type of sustainability relates to persistence and robustness of data infrastructures, legal and regulatory systems and knowledge communities – often in the context of very unstable political / economic situations. Barriers to data sharing are usually more cultural than technical. The data accessible to us is in itself a syntactic model of what mattered to the original funders (for example, REDD+ and carbon capture); this leads to data being reused for inappropriate purposes and increased uncertainty. 2 practical challenges from my domain: (1) Metadata and quality information in citizen science. (2) Reasoning about intersection between polygons representing protected areas and species ranges when these are bounded by lines of different types and of varying mobility (coastlines, political boundaries, physical fences) but the topological model necessary to capture this nuance is no longer in common use by commercial or open source GIS software.

3.2 Beyond Scientific Rationality: Why we need Critical Systems Thinking

Christoph Becker (University of Toronto, CA)

This talk outlined the normative character and inevitable value basis of modelling activities in computing and showed why and how critical perspectives are needed to address the key question of practical reason in modelling for sustainability: "How can we rationally justify the normative consequences of our models?".

Models in science and engineering are traditionally based on the tradition of the scientific method and inherit its realist ontology and objectivist epistemological foundations. Software

Engineering and Computing are based on this tradition, but are both about what *is* and what *ought to be*, i.e., design. The resulting models are enacted into behavior that acts with the world and changes it. Modelling involves assumptions that are in turn contingent upon unspoken beliefs – there are always decisions outside the model's justification. Underneath the surface lie moral and political decisions based on values. The talk offered a little Devil's Dictionary [1] of modelling to illustrate that computing tends to overlook these.

Unfortunately, scientific rationality misapplied to social systems [2] often fails to take into account the 'purposeful nature' of humans and social systems [3]. In assuming an objective goal function is given and unproblematic, it fails to account for the fact that it is often that goal function that is problematic [5], and that multiple contradictory views on the issues arising in a situation cannot be resolved away using scientific logic [4]. As a consequence, scientific rationality often suggests that the decisions it cannot reason about are simply irrational – and that is a mistake: This perspective "reduces practical reason to theoretical reason" [6] and ultimately fails to be relevant to the question at hand [7].

Since all models relevant to sustainability have normative consequences, they also have to be *legitimate*. The scientific method cannot legitimate them, because it has no access to values, moral and politics [8]. Neither can engineering theory on its own: Instrumental rationality cannot legitimate the normative implications of its own consequences [6], because it similarly cannot reason rationally about values, moral and politics. This makes it no less important to address the key question. Critical Theory and Critical Systems Thinking are essential perspectives that must be considered to begin addressing that challenge. To do so, computing must collaborate deeply with social disciplines. This raises well-known challenges [9], but cannot and must not be avoided.

References

1 Ambrose Bierce. *The devil's dictionary*. Wordsworth Editions, 1996
2 Herbert A. Simon. *The sciences of the artificial*. Cambridge, MA. 1969
3 Russell L. Ackoff. Towards a system of systems concepts, *Management Science*, 17(11), 1971
4 Horst WJ. Rittel and Melvin M. Webber. Dilemmas in a general theory of planning. *Policy Sciences*, 4(2):155-169, 1973
5 Peter Checkland. *Systems thinking, systems practice*. New York: J. Wiley, 1981
6 Werner Ulrich. *Critical heuristics of social planning: A new approach to practical philosophy*. Bern: P. Hapt., 1983
7 West C. Churchman.The artificiality of science. *Contemporary Psychology: A Journal of Reviews*. 15(6):385–386, 1970
8 West C. Churchman. *The systems approach and its enemies*. Basic Books, 1979
9 Geoffrey Bowker, Susan L. Star, William Turner, and Les Gasser (Eds.). *Social science, technical systems, and cooperative work: Beyond the great divide*. Mahwah, N.J.: Lawrence Erlbaum, 1997

3.3 Modelling with the Life Cycle Assessment (LCA) Framework

Didier Beloin-Saint-Pierre (Empa-Akademie – Zürich, CH)

The life cycle assessment (LCA) framework defines different modelling choices and assumptions to provide an evaluation of the environmental sustainability of products, services or systems.

The basis of this assessment is made from the quantitative comparison of environmental impacts from different options (i.e. products, services or systems) which have equivalent functions (e.g. provide electricity). These evaluated impacts cover a wide range of indicators (e.g. global warming potential, ecotoxicity, primary energy use) that occur all over the world and within the full life cycle of the considered options (i.e. natural resource extraction, manufacturing, use of product/service and end-of-life management). This comprehensive picture is essential to offer an assessment of sustainability and also helps in detecting potential impact displacement between indicators, regions and periods of time. When the modelling is finished, the option with lower environmental impacts is then considered more environmentally sustainable than the others [1]. This type of conclusion (i.e. more or less sustainable) highlights that LCA typically performs best when it offers a relative assessment of sustainability.

The environmental sustainability assessment that is carried out with the LCA framework uses two types of model. The first type describes the human activities with processes, product flows (i.e. link between processes) and their interactions with the environment (i.e. extracted natural resources and emissions of pollutants which are called elementary flows in LCA). The second type of model (i.e. life cycle impact assessment models) aggregates and translates the elementary flows into different environmental impacts. Such impacts can then be more or less aggregated into different indicators to offer information that fits the needs of decision-makers. Results can be presented with their uncertainties to provide more insights on the degree of confidence that LCA practitioners have on their work.

References

1 Stefanie Hellweg, and Canals M. Llorenc. Emerging approaches, challenges and opportunities in life cycle assessment. *Science*, 344(6188):1109–1113, 2014

3.4 The Role of Runtime Models for Decision Making in Sustainable Systems

Nelly Bencomo (Aston University – Birmingham, GB)

In this short presentation I talked about how to use models (design, runtime models and other models) to inform decision making wrt sustainability. I focus specially on the case of runtime models. I argue that models that support decision-making need to be updated over time. For example, new decisions will need to be reflected on the system while the loop continues. Decisions are related to the trade-offs between different quality properties related to sustainability. Runtime models can support the process of "IF-analysis" required to study the consequences of new decisions inserted.

3.5 Sustainability Debt: A Metaphor to Support Sustainability-Aware Software Systems Engineering

Stefanie Betz (KIT – Karlsruher Institut für Technologie, DE)

Joint work of Stefanie Betz, Christoph Becker, Ruzanna Chitchyan, Leticia Duboc, Steve M. Easterbrook, Birgit Penzenstadler, Norbert Seyff, Colin C. Venters
Main reference Stefanie Betz, Christoph Becker, Ruzanna Chitchyan, Leticia Duboc, Steve M. Easterbrook, Birgit Penzenstadler, Norbert Seyff, Colin C. Venters: "Sustainability Debt: A Metaphor to Support Sustainability Design Decisions", in Proc. of the Fourth International Workshop on Requirements Engineering for Sustainable Systems, RE4SuSy 2015, co-located with the 23rd IEEE International Requirements Engineering Conference (RE 2015), Ottawa, Canada, August 24, 2015., CEUR Workshop Proceedings, Vol. 1416, pp. 55–53, CEUR-WS.org, 2015.
URL http://ceur-ws.org/Vol-1416/Session2Paper4.pdf

This talk introduces the concept of sustainability debt. The metaphor helps in the discovery, documentation, and communication of sustainability issues in requirements engineering. Sustainability debt builds on the existing metaphor of technical debt and extend it to four other dimensions of sustainability to help think about sustainability-aware software systems engineering. It highlights the meaning of debt in each dimension and the relationships between those dimensions. Finally, it discusses the imitations and challenges of the metaphor sustainability debt.

3.6 Modelling for Natural Flood Management

Keith Beven (Lancaster University, GB)

Main reference Peter Metcalfe, Keith Beven, Barry Hankin, Rob Lamb: "A modelling framework for evaluation of the hydrological impacts of nature-based approaches to flood risk management, with application to in-channel interventions across a 29-km2 scale catchment in the United Kingdom", in Hydrological Processes, Vol. 31(9), pp. 1734–1748, 2017.
URL https://doi.org/10.1002/hyp.11140

Natural Flood Management (NFM) is proposed as a way of mitigating the damages of significant floods by distributed land management and storage of water in upstream catchment areas so as to reduce peak flows in areas at risk of flooding. This implies some decisions about investment in NFM measures, with assessment of resulting benefits (and potential dis-benefits). This will normally be achieved by modelling the impact of implementing NFM measures using hydrological runoff generation models that cascade inputs to hydraulic flood routing and flood inundation models. Impacts will often be assessed with reference to past historical floods, but because it is expected that flood frequencies might be changing as a result of climate change, might also involve some assessment of what future climate impacts on rainfalls and evapotranspiration might mean (using some form of weather generator model based on an ensemble of climate models) [1]. This represents a specific example of Modelling for sustainability for an environmental problem which involves significant sources of epistemic uncertainty, including the representation of runoff and flood routing processes; errors in input and evaluation data; effective values of model parameters; and potential scenarios of future boundary conditions.

References

1 Peter Metcalfe, Keith Beven, Barry Hankin, and Rob Lamb. A modelling framework for evaluation of the hydrological impacts of nature based approaches to flood risk management, with application to inchannel interventions across a 29km2 scale catchment in the United Kingdom. *Hydrological Processes*, 31(9):1734–1748, 2017

3.7 Working Together for Digitally Inspired Environmental Science

Gordon Blair (Lancaster University, GB)

Joint work of Ensemble Projects, Lancaster University
Main reference William Simm, Faiza Samreen, R. Bassett, Maria Angela Ferrario, Gordon S. Blair, Jonathan Whittle, P. J. Young: "SE in ES: opportunities for software engineering and cloud computing in environmental science", in Proc. of the 40th International Conference on Software Engineering: Software Engineering in Society, ICSE (SEIS) 2018, Gothenburg, Sweden, May 27 – June 03, 2018, pp. 61–70, ACM, 2018.
URL http://dx.doi.org/10.1145/3183428.3183430

Environmental science is at a crossroads as it is faced with new scientific challenges around climate change. There is a pressing need for a new kind of science in respond to this challenge, that is a science that is more open, integrated and collaborative. To support this, there is an equal need for new tools and techniques to support this style of science.

Ensemble is an umbrella initiative examining the role of technology in supporting this new kind of environmental science. It is a fundamentally trans-disciplinary initiative involving data scientists, computer scientists, experts in communication and also earth and environmental sciences. The broader programme is looking at a range of technologies, specifically new means of data acquisition at different scales (from the use of Internet of Things technology through to remote sensing), new techniques for making sense of the resultant rich but highly heterogeneous data sets (through emerging data science techniques tailored for the needs of environmental science), and also new technological infrastructure offering the elastic capacity for the storage and processing of this data (through the use of cloud computing).

A key aspect of this research is supporting environmental modelling in the cloud, with this work being carried out as part of the EPSRC-funded project "Models in the Cloud: Generative Software Frameworks to Support the Execution of Environmental Models in the Cloud" (EP/N027736/1). This work aims to make it easier for environmental modellers to run modelling experiments in the cloud through exploiting contemporary software engineering techniques, most notably model-driven engineering, to raise the level of abstraction of such platforms. This talk will explore results from this research project, including the application of such techniques in two contrasting areas of environmental science (related to climate science and hydrology respectively).

3.8 Modeling for Sustainability: the Software Engineering Perspective

Ruzanna Chitchyan (University of Bristol, GB)

In this talk I present the perspective that Software Engineering discipline has traditionally taken on the modelling and software generation activities. This traditional viewpoint is contrasted to the challenges posed by the Sustainability concern. In particular, I argue that:

(i) the constantly evolving notion of sustainability requires constant re-evaluation and adaptation of the models;

(ii) sustainability implies impact consideration at planetary scale, thus models of planetary scale are essential if the impact of the developed system is to be understood;

(iii) sustainability requires close consideration of social concerns, thus necessitating close, explicit, and continuous integration between social and technical models.

3.9 Modeling for Sustainability: Or How to Make Smart CPS Smarter?

Benoit Combemale (University of Toulouse, FR)

Main reference Benoît Combemale, Betty H. C. Cheng, Ana Moreira, Jean-Michel Bruel, Jeffrey G. Gray: "Modeling for sustainability", in Proc. of the 8th International Workshop on Modeling in Software Engineering, MiSE@ICSE 2016, Austin, Texas, USA, May 16-17, 2016, pp. 62–66, ACM, 2016.
URL http://dx.doi.org/10.1145/2896982.2896992

Various disciplines use models for different purposes. An engineering model, including a software engineering model, is often developed to guide the construction of a non-existent system. A scientific model is created to better understand an existing phenomenon (i.e., an already existing system or a physical phenomenon). An engineering model may incorporate scientific models to build a smart cyber-physical system (CPS) that require an understanding of the surrounding environment to decide of the relevant adaptation to apply. Sustainability systems, i.e., smart CPS managing resource production, transport and consumption for the sake of sustainability (e.g., smart grid, city, farming system), are typical examples of smart CPS. Due to the inherent complex nature of sustainability that must delicately balance trade-offs between social, environmental, and economic concerns, modeling challenges abound for both the scientific and engineering disciplines.

In this talk, I present a vision that promotes a unique approach combining engineering and scientific models to enable informed decision on the basis of open and scientific knowledge, a broader engagement of society for addressing sustainability concerns, and incorporate those decisions in the control loop of smart CPS. I introduce a research roadmap to support this vision that emphasizes the socio-technical benefits of modeling.

3.10 Modeling for Sustainability: How Quality Requirements Contribute to Sustainability?

Nelly Condori-Fernandez (Free University Amsterdam, NL)

Main reference Nelly Condori-Fernández, Patricia Lago: "Characterizing the contribution of quality requirements to software sustainability", Journal of Systems and Software, Vol. 137, pp. 289–305, 2018.
URL http://dx.doi.org/10.1016/j.jss.2017.12.005

The assessment and design based on the notion of sustainability requirements are still poorly understood. There is no consensus on which sustainability requirements should be considered. This talk introduces briefly the meaning about modeling for sustainability and highlights the importance of involving stakeholders for identifying requirements that can contribute the economic, technical, environmental and social sustainability dimensions of software-intensive systems. Also we argue that the relevance of the different dimensions depends on the type of software system.

With the purpose of defining a context-dependent sustainability model for software intensive systems, we present the design and main results of a survey that involves different target audiences (e.g. software architects, ICT practitioners with expertise in Sustainability, requirements engineers, and project managers).

3.11 Modelling Sustainability in Technology Transfer

Letícia Duboc (Ramon Llul University, ES)

Much of the software-based technology that surrounds our lives have their roots in universities research labs. However, transferring technology from the labs to the society is a complex process and putting in place a strategy to do so effectively is very challenging. This talk discusses some of the challenges on modelling sustainability in the context of technology transfer.

3.12 Modelling for Sustainability

Joao Goncalves (Empa-Akademie – Zürich, CH)

"Two examples and considerations about modelling"

The first presented modelling domain relates to microscopic traffic simulation. As a general consideration, modelling and simulation can be used not only as a means to visualise, interpret and quantify a system, but also as a replacement of field testing. In particular regard to sustainability, this substitution results in a compression of test times and can effectively reduce the necessary resources to conduct experiments.

The second presented application describes an ongoing work that attempts to discover and evaluate pathways to a post-fossil Switzerland. Translating the system outputs to

stakeholders and further applying their actions to the system constitutes an additional challenge.

Modelling for sustainability can be defined as the usage of modelling methodologies to consciously "optimise" sustainable usage of resources. However, if such studies are to have a significant impact on sustainability, Informatics must be used as a tool to translate complex and domain-specific assessments to non-experts and decision makers.

3.13 Modeling for Sustainability

Øystein Haugen (Ostfold University College – Halden, NO)

What is a Model? What is Modeling? What is Modeling for Sustainability? Models should execute/behave and mimic a referent system. Modeling is the creation and evolution of a model. Modeling for sustainability is when the referent system is concerned with sustainability.

3.14 Modeling to Reduce Waste in Chemical Production

Øystein Haugen (Ostfold University College – Halden, NO) and Per-Olav Hansen

This talk presented a use-case in the European ECSEL project Productive4.0 owned by Unger Fabrikker in Norway and executed by the Norwegian consortium in Productive4.0. The use-case is about reducing the waste originating from the transition period between the production of two high-quality chemical products. During the transition period, no proper product is produced and this produce must be further handled as waste. The purpose of the use-case is to find models that can make it possible to reduce the transition period and the amount of potential waste.

3.15 Reflections on Marvin Minsky's Definition of "Model"

Lorenz Hilty (Universität Zürich, CH)

"To an observer B, an object A* is a model of an object A to the extent that B can use A* to answer questions that interest him about A." [1], p. 426.

This definition includes concrete, tangible models as well as conceptual models that are described in (usually formal) languages. The definition has some fruitful implications:

- The purpose of a model can be described by specifying the type of questions the model is intended to answer about the original (A).
- The purpose of a model is connected to the (epistemic) interest of an observer.

- There can be multiple models of the same original, depending on the purpose. Zeigler calls two models complementary if they embody consistent hypotheses about the original (but in a different way), and competitive if they embody mutually exclusive hypotheses [2], p. 13.
- The term "representation" can be avoided in the definition of "model". I consider this an advantage because "representation" is a term that raises many epistemological issues.
- The terms "abstraction" and "simplification" can be avoided in the definition of "model". Characterizing models as abstractions or simplifications implies there could be an entity that is "no abstraction" or "no simplification" of another entity (something like a "perfect copy", which is however not a model because it is no abstraction or simplification), an idea that again raises epistemological issues.
- Both descriptive and prescriptive models can be subsumed under Minsky's definition, namely in the following way: If A already exists, the model is descriptive, otherwise (if A is to be created), the model is prescriptive. In the latter case, the "questions that interest [the observer] about A" are addressing the consequences of design decisions regarding A. If A exists and we are interested in future changes A may undergo (or of impacts A will be subjected to), the model is descriptive and prescriptive.
- We can ask how a model is used to generate answers to questions. "Generating aswers" is often done by some sort of experimentation, e.g., by setting parameters and initial conditions to create an instance of the model and let an algorithm interpret it. Simulation can thus be defined as experimenting with a model. (If there is no need for experimentation, we are in the exceptional situation that the model is simple enough to be treated analytically.)
- Because a model is not a statement, but a generator for a (usually infinite) set of statements about the original, it is usually not verifiable, but falsifiable.

References

1 Marvin L. Minsky. *Matter, Mind and Models*. MIT Press, Boston MA, USA, 1968
2 Bernard P. Zeigler. *Multifacetted Modelling and Discrete Event Simulation*. Academic Press, London, GB, 1984

3.16 Sustainability: Scientific Theories and Models

Jean-Marc Jézéquel (IRISA – Rennes, FR)

Joint work of Diverse Team, at IRISA/Inria Rennes

A Model is an abstraction of an aspect of the world for a specific purpose. Therefore a Scientific Theory for supporting sustainability is a Model (but a Model is not always a Scientific Theory, eg. because it might not be falsifiable). In facts, Creating a Scientific Theory is (evermore) Writing Software. Conversely writing (useful) Software is like Creating a Scientific Theory, with validation tests playing the role of experiments in science. In this talk we explore how MDE technology could be used to support scientific theories of sustainability.

3.17 Modeling of Sustainability: Sustainable Software Engineering

Eva Kern (Universität Lüneburg, DE)

Main reference Stefan Naumann, Markus Dick, Eva Kern, Timo Johann: "The GREENSOFT Model: A Reference Model for Green and Sustainable Software and its Engineering," In: Sustainable Computing: Informatics and Systems, 1(4):294–304, 2011.
URL https://doi.org/10.1016/j.suscom.2011.06.004

This talk summarized the way of modeling for sustainability of software. Giving a general short introduction into models of the field of software engineering, it presents different kinds of models of the field of green and sustainable software engineering: a life cycle model for sustainable software products, a reference model for green software and its engineering, procedure models for green software engineering, a measurement model to analyze the consumption of energy and resources while using software, and a quality model for sustainable software.

References

1 Stefan Naumann, Markus Dick, Eva Kern, and Timo Johann. The greensoft model: A reference model for green and sustainable software and its engineering. *Sustainable Computing: Informatics and Systems*, 1(4):294-304 Elsevier, 2011
2 Markus Dick and Stefan Naumann. Enhancing Software Engineering Processes towards Sustainable Software Product Design. *EnviroInfo 2010. Integration of Environmental Information in Europe. Proceedings of the 24th International Conference on Informatics for Environmental Protection*, 706-715, 2010
3 Eva Kern, Markus Dick, Stefan Naumann, Achim Guldner, and Timo Johann. Green software and green software engineering–definitions, measurements, and quality aspects. *ICT for Sustainability. Proceedings of the First International Conference on Information and Communication Technologies for Sustainability*, 87-94, 2013
4 Markus Dick, Eva Kern, Jakob Drangmeister, Stefan Naumann, and Timo Johann. Measurement and rating of software-induced energy consumption of desktop PCs and servers. Pillmann, Werner; Schade, Sven; Smits, Paul (Eds.): *Innovations in Sharing Environmental Observation and Information. Proceedings of the 25th EnviroInfo Conference "Environmental Informatics*, 290-299, 2011

3.18 Modelling for Sustainability in the Now

Jörg Kienzle (McGill University – Montreal, CA)

This short presentation takes a critical look at our ever increasing capability to model and make predictions about the future. The talk points out the very real possibility of abuse of this capability if it is not made available to the population at large. The talk ends by pointing out the inevitability of change, and consequently the continuous need to adapt.

3.19 Modeling for Sustainability: Challanges and Modeling Examples in Green Software

Sedef Akinli Kocak (Ryerson University – Toronto, CA)

Joint work of Sedef Akinli Kocak, Gulfem Alptekin, Ayse Basar Bener, Patricia Lago, Ivica Crnkovic, Birgit Penzenstadler

Different models have been developed and used in science and engineering disciplines. This talk gives a short introduction into what modeling is in general and summarizes main purpose of modeling for sustainability and challenging issues. The main challenging issues include taking interdisciplinary approach, managing uncertainty, taking long-term and global-local perspectives, and stakeholders participation with integration of their values and objectives. Then different modeling efforts have been presented in the area of green software and software for sustainability. The first presented one is based on modeling energy consumption of software products based on quantitative analysis [1]. The second presented modeling effort is multi-criteria decision making for software quality model. The third presented modeling is framing sustainability as a product quality [2]. The talk finalizes with take away question: how can models be developed and/or improved for sustainability purposes and used in support of decision-making?

References

1 Sedef Akinli Kocak, Gulfem Alptekin, and Ayse Basar Bener. Integrating environmental ssustainability in software product quality. in *Proceedings of the 4th International Workshop on Requirements Engineering for Sustainable Systems co-located with RE 2015*, 1416(0074-1416-7):17-24. 2015

2 Patricia Lago, Sedef Akinli Kocak, Ivica Crnkovic, and Birgit Penzenstadler. Framing sustainability as a property of software quality. *Communications of the ACM*, 58(4):70-78, 2015

3.20 Models of Programming Languages

Peter D. Mosses (TU Delft, NL)

Joint work of L. Thomas van Binsbergen, Neil Sculthorpe, Peter D. Mosses, et al.
Main reference L. Thomas van Binsbergen, Neil Sculthorpe, Peter D. Mosses: "Tool support for component-based semantics", in Proc. of the Companion 15th International Conference on Modularity, Málaga, Spain, March 14–18, 2016, pp. 8–11, ACM, 2016.
URL http://dx.doi.org/10.1145/2892664.2893464

My research domain is meta-languages and tool support for specifying models of programming languages. This brief presentation first recalled the main features of such models, and the kinds of meta-languages typically used to specify them.

The PLanCompS project has developed a component-based approach to modelling programming languages. The semantics of each language construct is specified by translating it to an open-ended library of so-called 'funcons' (fundamental programming constructs). The behaviour of each funcon is fixed, and its definition does not change when new funcons are added. The beta-release of an initial library of funcons is available, together with some illustrative component-based specifications (at https://plancomps.github.io/CBS-beta).

The component-based approach supports reuse and co-evolution when modelling programming languages. This could encourage use of formal models by language developers, which might lead to better language design, and perhaps ultimately reduce waste of resources due to software bugs and lack of portability, but there appears to be no direct relevance to modelling for sustainability. It might however be interesting to investigate whether the component-based approach could be exploited for general modelling.

3.21 MDE and Sustainability: Questions

Gunter Mussbacher (McGill University – Montreal, CA)

Model-Driven Engineering (MDE) is based on the premise that each model conforms to a well-defined language (metamodel / grammar / profile) which specifies the concepts and relationships of a domain (i.e., abstract syntax), their representation (i.e., concrete grammar), as well as their meaning (i.e., semantics). A model is an abstraction of reality. It is a simplified, purposeful representation of a specific property/quality adjusted to human needs, hence reducing complexity to the human scale. To be useful, a model must be accurate and concise. Given these characteristics, a model enables humans to understand a domain, to communicate, to reason about it and make predictions about the property/quality, and – in some cases – implement the system. Is sustainability just another quality that can be handled like all other qualities? What is different? Is it the heterogeneity of the set of required models? Is it the uncertainty that needs to be reflected in the models? Is it continuous systems vs. discrete systems? Is it system thinking vs. divide and conquer?

3.22 Modeling for Sustainability in Software Engineering

Oscar M. Nierstrasz (Universität Bern, CH)

Software Engineering (SE) models abstract from objects of a given domain, whether real or virtual, in order to support reasoning or communication. Models may be descriptive, describing existing artifacts, or prescriptive, specifying something yet to be built. Models may support sustainability of SE processes, i.e., to ensure sustainable cost over time, or sustainability of SE artifacts, i.e., to ensure that code will be maintainable in the long term.

3.23 Modeling and Sustainability: Fitness-for-Purpose and Process

Richard F. Paige (University of York, GB)

Models are created for a purpose, and must ultimately be judged as fit for that purpose. What measures, metrics and qualities are important for understanding fitness-for-purpose for sustainability models? In traditional software engineering we are concerned with qualities

such as correctness and consistency, whereas for sustainability (where models may live on for many years and may be managed by different people with different skills) other qualities, such as habitability, may be more important.

Similarly, models are created following different processes, including bottom-up (based on examples), top-down (using a domain-specific modeling language or general-purpose modeling language), via democratic process, via automatic generation, etc. Some models are left partially tacit or implicit via certain modeling processes. What is a suitable process for engineering the complex heterogeneous modeling collections that are needed for sustainability engineering?

3.24 5 Dimensions of Sustainability, Sustainability Analysis Diagram, and Leverage Points

Birgit Penzenstadler (California State University, US)

Joint work of Christoph Becker, Stefanie Betz, Ruzanna Chitchyan, Leticia Duboc, Steve M. Easterbrook, Birgit Penzenstadler, Norbert Seyff, Colin C. Venters
Main reference Christoph Becker, Stefanie Betz, Ruzanna Chitchyan, Leticia Duboc, Steve M. Easterbrook, Birgit Penzenstadler, Norbert Seyff, Colin C. Venters: "Requirements: The Key to Sustainability", IEEE Software, Vol. 33(1), pp. 56–65, 2016.
URL http://dx.doi.org/10.1109/MS.2015.158

In the area of requirements engineering we, inter alia, use models to illustrate concepts and come to agreements about the context and the system under development amongst a wide range of stakeholders. For that, we use five dimensions of sustainability (individual, social, economic, technical, and environmental) as well as three orders of effects (immediate, enabling, and structural) and depict a summary of this in a sustainability analysis diagram. Furthermore, we have applied the concept of leverage points (cf. Donella Meadows) to software systems for sustainability to understand how developers can use systems thinking to consider their designs in a larger context.

- What is your definition of modeling? Modeling is the abstraction from and representation of real world to conceptual elements and relations between them. In requirements engineering, we do this in often informal or semi-formal and illustrative ways.
- What is meant by "modeling for sustainability" in your domain/area of work? We try to develop (software) systems that support the use of our planet preserving its capacity to support living on it.

3.25 Contributions in Software Engineering and Green IT

Lionel Seinturier (Lille I University, FR)

Main reference Maxime Colmant, Romain Rouvoy, Mascha Kurpicz, Anita Sobe, Pascal Felber, Lionel Seinturier: "The next 700 CPU power models", Journal of Systems and Software, Vol. 144, pp. 382–396, 2018.
URL http://dx.doi.org/10.1016/j.jss.2018.07.001

This talk summarizes my research expertises in sustainable computing, and the contributions made in two key domains in relation with the workshop: software engineering and green IT.

In terms of software engineering, I have some contributions in the model-driven engineering and self-adaptive systems communities. Some of my recent work especially deals with domain-specific language design and formal methods for specifying the reconfiguration space and the legal states a software system can be in. With my co-authors we have then applied some solutions coming from the control theory domain to generate some discrete event controllers that ensure that the system under control stays within the boundaries that have been specified. This work have been applied to the znn.com exemplar well-known in the self-adaptive system community.

In terms of green IT, me and my group of colleagues have developed in the recent years the PowerAPI (http://www.powerapi.org) library that enables to implement software-defined power meters to measure the energy induced by software systems. Among the goals that are pursued, we want to be able to identify energy hotspots in software systems, be able to rank web sites and services according to their energy footprint, and infer the energy model of hardware components. On this last point, we especially showed that the heterogeneity of modern CPU is so vast that one cannot a priori define a realistic power model (due to some very high variability, and a very high number of available hardware performance counters to potentially monitor). To solve this problem, we devised a solution where we applied some machine learning techniques to learn the power models.

3.26 Modeling for Sustainability: Lessons from Air Quality Decision-Making

Noelle Selin (MIT – Cambridge, US)

Sustainability is a critical challenge for engineering research and education as a (and perhaps the) critical societal challenge of the 21st century. For sustainability science, key core questions involve modeling: specifically, how theory and models can be formulated to better account for human-environment interactions, and how society can most effectively guide human-environment systems towards a sustainability transition [1]. Previous research suggests lessons for researchers (and modelers) interested in crafting usable knowledge for sustainable development [2]. Synthesis of previous scientific assessment efforts has shown that for research to be effective in influencing policy, it needs to be perceived by stakeholders to be credible, salient, and legitimate [3].

In the domain of air quality, useful lessons can be gleaned through efforts to understand the pathways from policies that control human activities and emissions, through the fate and transport of atmospheric pollutants, to exposure and health impacts. Simulating these pathways involves linking different sorts of models (economic, atmospheric, and health impact modeling) as well as accounting for system interactions and decision-making through case studies and policy experiments. In this talk, I address how the goal of having impact on sustainability-relevant societal challenges such as air quality can influence model-based research, through three mechanisms: scientific assessment processes, co-production with stakeholders, and co-production with boundary organizations. Examples of these different mechanisms in practice are described through examples of modeling mercury pollution [4], evaluating the impact of climate action on air quality outcomes [5], and assessing the co-benefits of U.S. climate policy [6]. I then examine how a policy-driven orientation can affect

decisions about model scale (setting boundaries and resolution [7], complexity (simplifying processes [8], and uncertainty evaluation (evaluating end-member analyses [9] and conducting model ensembles). Experiments in model-based decision-making are summarized, showing that when users engage themselves with models, individuals are more likely to find win-win sustainability trade-offs [10], and groups find consensus faster [11]. Sustainability effects of policies can also be evaluated quantitatively using frameworks from inclusive wealth accounting, as shown by a case study of non-fossil energy investment in Saudi Arabia [12]. Simple model equations can potentially be more effective than complex models in informing policy, such as new metrics to inform global mercury negotiations [13]. Further case studies are needed to help inform the development of new models and frameworks to address sustainability as a systems problem.

References

1 Kates, Robert W. *What kind of a science is sustainability science?*. Proceedings of the National Academy of Sciences 108.49: 19449-19450, 2011
2 Clark, William C., et al. *Crafting usable knowledge for sustainable development.* Proceedings of the National Academy of Sciences 113.17: 4570-4578, 2016
3 Cash, David W., et al. *Knowledge systems for sustainable development.* Proceedings of the National Academy of Sciences 100.14: 8086-8091, 2003
4 Selin, Noelle E., and Daniel J. Jacob. Seasonal and spatial patterns of mercury wet deposition in the United States: Constraints on the contribution from North American anthropogenic sources. *Atmospheric Environmen* 42(21):5193-5204, 2008
5 Garcia-Menendez, Fernando, et al. US air quality and health benefits from avoided climate change under greenhouse gas mitigation. *Environmental Science & Technology*, 49(13):7580-7588, 2015
6 Thompson, Tammy M., et al. A systems approach to evaluating the air quality co-benefits of US carbon policies. *Nature Climate Change*, 4(10):917, 2014
7 Thompson, Tammy M., and Noelle E. Selin. Influence of air quality model resolution on uncertainty associated with health impacts. *Atmospheric Chemistry and Physics*, 12(20):9753-9762, 2012
8 Brown-Steiner, Benjamin, et al. *Evaluating Simplified Chemical Mechanisms within CESM Version 1.2 CAM-chem (CAM4): MOZART-4 vs. Reduced Hydrocarbon vs. Super-Fast Chemistry*, Geosci. Model Dev. Discuss, 2018
9 Giang, Amanda, and Noelle E. Selin. Benefits of mercury controls for the United States. *Proceedings of the National Academy of Sciences*, 113(2):286-291, 2016
10 Czaika, Ellen, and Noelle E. Selin. Model use in sustainability policy making: An experimental study. *Environmental Modelling & Software*, 98:54-62, 2017
11 Czaika, Ellen, and Noelle E. Selin. *Taking action to reduce waste: quantifying impacts of model use in a multiorganizational sustainability negotiation. Negotiation and Conflict Management Research*, 9(3):237-255, 2016
12 Collins, Ross D., et al. Using inclusive wealth for policy evaluation: Application to electricity infrastructure planning in oil-exporting countries. *Ecological Economics*, 133:23-34, 2017
13 Selin, Noelle E. A proposed global metric to aid mercury pollution policy. *Science*, 360(6389):607-609, 2018

3.27 Requirements Engineering for Evolution Towards Sustainability

Norbert Seyff (FH Nordwestschweiz, CH)

Continuous requirements elicitation is an essential aspect of software product evolution to keep systems aligned with changing user needs. However, current requirements engineering approaches do not explicitly address sustainability in the evolution of systems. Reasons include a lack of awareness and a lack of shared understanding of the concept of sustainability in the RE community. Identifying and analysing the effects of requirements regarding sustainability is challenging, as these effects can have an impact on multiple stakeholders and manifest themselves in one or more sustainability dimensions at different points in time. We argue that tailored requirements engineering approaches are needed which allow the engagement of a large number of stakeholders (including users and domain experts) in a continuous cycle of negotiation regarding the potential effects of requirements on sustainability.

3.28 Software Architecture Modeling for Sustainability: WTFs/Minute

Colin Venters (University of Leeds, GB), Christoph Becker (University of Toronto, CA), Stefanie Betz (KIT – Karlsruher Institut für Technologie, DE), and Birgit Penzenstadler (California State University, US)

Joint work of Colin C. Venters, Rafael Capilla, Stefanie Betz, Birgit Penzenstadler, Tom Crick, Steve Crouch, Elisa Yumi Nakagawa, Christoph Becker, Carlos Carrillo
Main reference Colin C. Venters, Rafael Capilla, Stefanie Betz, Birgit Penzenstadler, Tom Crick, Steve Crouch, Elisa Yumi Nakagawa, Christoph Becker, Carlos Carrillo: "Software sustainability: Research and practice from a software architecture viewpoint", Journal of Systems and Software, Vol. 138, pp. 174–188, 2018.
URL http://dx.doi.org/10.1016/j.jss.2017.12.026

This talk highlights the role that software architectures play in the development of technically sustainable software. It is argued that software architectures are the primary carrier of system qualities (NFR) i.e. pre-system understanding, and influence how developers are able to understand, analyze, extend, test and maintain a software system i.e. post-deployment system understanding. As such, software architectures provide a mechanism for reasoning about quality attributes. This presentation proposes that sustainable software architectures are fundamental to the development of technically sustainable software to address architectural drift and erosion, decay, and architectural knowledge vaporization [1].

References

1 Colin C. Venters, Rafael Capilla, Stefanie Betz, Birgit Penzenstadler, Tom Crick, Steve Crouch, Elisa Yumi Nakagawa, Christoph Becker, and Carlos Carrillo. Software Sustainability: research and practice from a software architecture viewpoint. *Journal of Systems and Software*, 138:174-188, 2017

3.29 Human Values in Software Engineering – Where Are They?

Jon Whittle (Monash University – Clayton, AU)

After decades of research and practice in software engineering, a range of well-established methodologies have been developed that (generally speaking) help to produce software that has the right functionality, at an affordable cost, is safe, secure, and safeguards data privacy. However, there are a whole range of 'human values' that have not been considered in software engineering, such as gender diversity, transparency, integrity, social responsibility, family or corporate values. No software is values-free, however. And so this talk argues that software designers ought to explicitly consider human values in software design. We see this as a new paradigm in software engineering which has a number of challenges including how to specify human values, how to trace human values throughout the software lifecycle, and how to measure values in software.

3.30 Modelling for Sustainability – 5 minute introduction

Paul Young (Lancaster University, GB)

This presentation addresses the initial questions set for the workshop (what is modelling? what is modelling for sustainability?) from the point of view of an atmospheric/climate scientist. Our definition of modelling mirrors that used by several other fields represented here, but our ideas of "sustainability" are more coupled to global sustainability/sustainable development issues as defined by the 1987 Bruntland Report (https://en.wikipedia.org/wiki/Brundtland_Commission). Other relevant issues of sustainability relate to the sustainability of the code: well documented and re-usable, but also efficient from the point of view of energy consumption.

4 Working Groups

4.1 Modeling for Sustainability: A Multidisciplinary Problem

Betty H. C. Cheng (Michigan State University – East Lansing, US)

In addition to providing definitions for model and modeling for sustainability, this talk briefly overviews previous experiences of modeling for sustainability. Lessons learned and challenges from a sustainability project from 20 years earlier are interestingly still applicable in current day efforts with sustainability. The main difference is the scale and complexity of the models, data, and integration have increased dramatically, largely due to the technological advances in the past two decades. Three key challenges are highlighted: data access and integration; model integration; and role and impact of uncertainty.

References

1 Betty H. C. Cheng, Robert H. Bourdeau, and Gerald C. Gannod: The object-oriented development of a distributed multimedia environmental information system. *SEKE 1994*: 70-77, 1994

2 Joseph Sharnowski, Gerald C. Gannod, and Betty H. C. Cheng. A distributed, multimedia environmental information system. in *Proc. of IEEE Int. 32 Chapter 1 Conference on Multimedia and Computing Systems*, 1995

3 Betty H. C. Cheng, Robert H. Bourdeau, and BC Pijanowski. A decision support system for regional environmental analysis. in *Proceedings of the 25th Int'l Symp. on Remote Sensing and Global Environmental Change: Tools for Sustainable Development*, 2:223-233, 1993

4 Betty H. C. Cheng, Robert H. Bourdeau, and Bryan C. Pijanowski. A regional information system for environmental data analysis (with). *Journal of Photogrammetric Engineering and Remote Sensing*, 62(7):855–861, 1996

5 Benoit Combemale, Betty H. C. Cheng, Ana Moreira, Jean-Michel Bruel, and Jeff Gray. Modeling for Sustainability. in *Proceedings of the 8th International Workshop on Modeling in Software Engineering (MiSE '16)*. ACM, New York, NY, USA, 62-66. 2016.

Participants

- Olivier Barais
INRIA – Rennes, FR
- Lucy Bastin
Aston University – Birmingham, GB
- Christoph Becker
University of Toronto, CA
- Didier Beloin-Saint-Pierre
Empa-Akademie – Zürich, CH
- Nelly Bencomo
Aston University – Birmingham, GB
- Stefanie Betz
KIT – Karlsruher Institut für Technologie, DE
- Keith Beven
Lancaster University, GB
- Gordon Blair
Lancaster University, GB
- Gael Blondelle
Eclipse Foundation Europe GmbH – Zwingenberg, DE
- Betty H. C. Cheng
Michigan State University – East Lansing, US
- Ruzanna Chitchyan
University of Bristol, GB
- Benoit Combemale
University of Toulouse, FR
- Nelly Condori-Fernandez
Free University Amsterdam, NL
- Letícia Duboc
Ramon Llul University, ES
- François Fouquet
University of Luxembourg, LU
- Joao Goncalves
Empa-Akademie – Zürich, CH
- Øystein Haugen
Ostfold University College – Halden, NO
- Lorenz Hilty
Universität Zürich, CH
- Jean-Marc Jézéquel
IRISA – Rennes, FR
- Eva Kern
Universität Lüneburg, DE
- Jörg Kienzle
McGill University – Montreal, CA
- Sedef Akinli Kocak
Ryerson University – Toronto, CA
- Peter D. Mosses
TU Delft, NL
- Gunter Mussbacher
McGill University – Montreal, CA
- Oscar M. Nierstrasz
Universität Bern, CH
- Richard F. Paige
University of York, GB
- Birgit Penzenstadler
California State University, US
- Bernhard Rumpe
RWTH Aachen, DE
- Lionel Seinturier
Lille I University, FR
- Noelle Selin
MIT – Cambridge, US
- Norbert Seyff
FH Nordwestschweiz, CH
- Eugene Syriani
Université de Montréal – Quebec, CA
- Colin Venters
University of Leeds, GB
- Jon Whittle
Monash University – Clayton, AU
- Paul Young
Lancaster University, GB