DAGSTUHL
REPORTS

Volume 7, Issue 5, May 2017

Theory and Applications of Hashing (Dagstuhl Seminar 17181)
Martin Dietzfelbinger, Michael Mitzenmacher, Rasmus Pagh, David P. Woodruff, and Martin Aumüller 1

Theory of Randomized Optimization Heuristics (Dagstuhl Seminar 17191)
Carola Doerr, Christian Igel, Lothar Thiele, and Xin Yao 22

Human-Like Neural-Symbolic Computing (Dagstuhl Seminar 17192)
Tarek R. Besold, Artur d'Avila Garcez, and Luis C. Lamb 56

Formal Synthesis of Cyber-Physical Systems (Dagstuhl Seminar 17201)
Calin A. Belta, Rupak Majumdar, Majid Zamani, and Matthias Rungger 84

Challenges and Opportunities of User-Level File Systems for HPC (Dagstuhl Seminar 17202)
André Brinkmann, Kathryn Mohror, and Weikuan Yu 97

Geometric Modelling, Interoperability and New Challenges (Dagstuhl Seminar 17221)
Falai Chen, Tor Dokken, Thomas A. Grandine, and Géraldine Morin 140

Robust Performance in Database Query Processing (Dagstuhl Seminar 17222)
Renata Borovica-Gajic, Goetz Graefe, and Allison Lee 169

Dagstuhl Reports, Vol. 7, Issue 5 ISSN 2192-5283

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
January, 2018

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.7.5.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 17181

Theory and Applications of Hashing

Edited by

Martin Dietzfelbinger[1], Michael Mitzenmacher[2], Rasmus Pagh[3], David P. Woodruff[4], and Martin Aumüller[5]

1 **TU Ilmenau, DE,** martin.dietzfelbinger@tu-ilmenau.de
2 **Harvard University – Cambridge, US,** michaelm@eecs.harvard.edu
3 **IT University of Copenhagen, DK,** pagh@itu.dk
4 **IBM Almaden Center – San Jose, US,** dpwoodru@us.ibm.com
5 **IT University of Copenhagen, DK,** maau@itu.dk

Abstract

This report documents the program and the topics discussed of the 4-day Dagstuhl Seminar 17181 "Theory and Applications of Hashing", which took place May 1–5, 2017. Four long and eighteen short talks covered a wide and diverse range of topics within the theme of the workshop. The program left sufficient space for informal discussions among the 40 participants.

Seminar May 1–5, 2017 – http://www.dagstuhl.de/17181

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity, H.3 Information Storage and Retrieval

Keywords and phrases connections to complexity theory, data streaming applications, hash function construction and analysis, hashing primitives, information retrieval applications, locality-sensitive hashing, machine learning applications

Digital Object Identifier 10.4230/DagRep.7.5.1

1 Executive Summary

Martin Dietzfelbinger
Michael Mitzenmacher
Rasmus Pagh
David P. Woodruff
Martin Aumüller

Background

The idea of hashing was proposed in the 1950s as an efficient method for implementing symbol tables in compilers. In succeeding decades it has emerged as an algorithmic tool that goes well beyond its original purpose, providing solutions for a wide range of algorithmic problems. While the theory of hashing may appear mature, in fact, many new advances have been made in recent years. Also, the number of applications has grown to an extent that few people realize how broad the reach of hashing is, or have a comprehensive overview. The aim of this seminar was to bring together researchers with an interest in hashing methods (and more generally random mappings) from various perspectives, spanning theory and a diverse set of application areas. In this way we wanted to identify opportunities for further advancing the field.

Theory

The theoretical side includes aspects such as the design of hash functions, the design of algorithm and data structure primitives, and the mathematical analysis of hashing schemes. For hash function design, since Carter and Wegman proposed universal hashing there has been a fertile research agenda identifying sufficient randomness properties of hash functions for various applications, and on the other hand devising resource-efficient hash functions having these properties. While new simple and efficient hash function constructions with strong theoretical properties keep appearing (seminar participants have contributed to this), there are still many applications of hashing for which no efficient hash function construction is known. At the same time it is increasingly clear that traditional measures of hash function performance like collision probability and independence are inadequate to describe the randomness properties needed in many applications.

While hashing is interesting in its own right, it has also become a foundational building block for higher level algorithms and data structures, each of which can in turn often find uses in a variety of application spaces. Various hash-based sketches laid the groundwork for the field of streaming algorithms, by showing how approximate counting could be done effectively. More recently, hashing algorithms have provided frameworks for similarity measures for text and images, data reconciliation, and even fast sparse Fast Fourier transform algorithms. As we construct richer, more elaborate structures on top of the hashing building blocks, we stretch what we require from hashing further.

Mathematical analysis of hashing schemes, an area enriched by early work of Knuth but grown well beyond, has inspired the development of a number of combinatorial methods and results. In the 1990s it was realized that load balancing using the best of two (or more) random choices leads to a much more even distribution, often referred to as the "power of two choices". Another great success in this area, obtained during the last decade, has been the analysis of cuckoo hashing; several seminar participants were involved in this. On the other hand, as questions become answered, new questions arise; in the case of cuckoo hashing, the effectiveness of variants including double hashing and partial-key cuckoo hashing are not understood. Beyond the hashing schemes themselves, data structures and algorithms that make use of lower-level hashing primitives further require and benefit from mathematical analysis. The fundamental connections between hashing, sparse reconstruction algorithms, and various sketching approaches are only now starting to be realized.

Applications

Hashing is of course heavily used in information storage and retrieval contexts. For example, the "power of two choices" paradigm(where several seminar participants were among the pioneers) has resulted in extremely scalable and robust architectures for distributed hash tables (also known as key-value stores).

Other applications of hashing are appearing at a tremendous rate, as systems-designers and builders become accustomed to a world where approximate answers (as opposed to exact answers) are not only sufficient, they are necessary for efficiency. Indeed, hashing was one of the key methodologies for handling big data well before "big data" was even a widely used term. Since the seminal paper of Flajolet and Martin that showed how to efficiently compute an approximate count of the number of elements in a data stream, hashing has been a central tool in the design of algorithms for data streams, where the inability to store all the information that passes requires approximations. But in recent years the use of hashing has spread to many other settings where data is stored and accessible, but scalability can be achieved only by resorting to approximation algorithms similar to those developed in the setting of data streams. One early success story in this direction is the method for identifying near-duplicate web pages in Altavista using min-wise hashing. Another is HyperANF, a refined version of the Flajolet-Martin method that was used in 2012 to compute the distance distribution of the Facebook social network, making it by far the largest scale Milgram-like experiment ever performed. Finally, Bloom filters, invented around 1970 to provide a small-memory approximate representation of a set, have become a staple in systems, with countless variations expanding on its initial functionality, such as counts associated with set elements or aging out of set items as new elements are dynamically added.

In the field of machine learning, random mappings of data to lower-dimensional vectors that are easier to handle is of increasing importance for big data applications. This is true in particular since machine learning algorithms often work with kernelized feature vectors whose dimension far exceeds the size of data vectors. Designing randomized mappings that meet the criteria of machine learning applications has been an active research area in recent years. NIPS 2014 awarded one of two best paper awards a paper co-authored by seminar participant Shrivastava that presents a new asymmetric locality-sensitive hash function design for machine learning applications and shows how it leads to significant speedups.

The rich interplay between theory and practice in the field of hashing cannot be overstated. Applications drive the need for new algorithms and data structures based on hashing, as well as the need for more efficient classes of hash function families with provable theoretical guarantees. Specific implementations developed in the field often do not have theoretical guarantees on performance or accuracy, creating new theoretical problems and driving a need to make theory and practice meet.

Industrial relevance. The workshop topic was highly relevant for companies dealing with big data. Three seminar participants are affiliated with Google, one has worked on hashing algorithms at AT&T for a decade, and one is affiliated with VMware. Also, one organizer was affiliated with IBM at the time of the seminar.

Outcome of the seminar

The seminar brought together quite a few leading figures in the area of hashing, mixed with a good fraction of young researchers, some of which are behind many of the most exciting results in the area in recent years. Areas that were particularly well represented were: Analysis of multiple-choice hashing methods, hashing for high-dimensional search problems, applications of hashing in string algorithms, applications of hashing in machine learning, streaming (approximation) algorithms, high-performance hash function construction, and algorithm engineering. Many results in these areas were presented in 18 shorter talks. Four longer talks (by A. Andoni, A. McGregor, U. Wieder, and Q. Zhang) contributed background, overview over new results, and aspects of applications of hashing in industry. Open problems were discussed in an open problem session; four of them are included in this report.

The paper [1] on fillable arrays co-authored by J. Nelson was motivated by a talk given by T. Hagerup (see Section 3.7) and can thus be seen as a first direct result of the seminar.

We, the organizers, would like to thank all participants for their contributions in talks, material, and discussions, and in particular the speakers of the longer talks. Many thanks are due to the Dagstuhl staff both in the offices and in the castle for their support in making this seminar a success.

References

1 Jacob Teo Por Loong, Jelani Nelson, Huacheng Yu: Fillable arrays with constant time operations and a single bit of redundancy. CoRR abs/1709.09574 (2017)

2 Table of Contents

Executive Summary

Martin Dietzfelbinger, Michael Mitzenmacher, Rasmus Pagh, David P. Woodruff, and Martin Aumüller . . . 1

Overview of Talks

Beyond Locality Sensitive Hashing
Alexandr Andoni . . . 6

Distance-Sensitive Hashing
Martin Aumüller . . . 6

Applications of Hashing in Semantic Search
Hannah Bast . . . 7

Set Similarity Search Beyond MinHash
Tobias Christiani . . . 7

BDDs for Minimal Perfect Hashing: Merging Two State-Space Compression Techniques
Stefan Edelkamp . . . 8

Synchronization Strings: Near-Optimal Codes for Insertions and Deletions
Bernhard Haeupler . . . 8

On-the-Fly Array Initialization in Less Space
Torben Hagerup . . . 9

An Adaptive Sublinear Time Block Sparse Fourier Transform
Michael Kapralov . . . 9

Locality Sensitive Distortion
Ravi Kumar . . . 10

Improved ℓ_2/ℓ_2 Sparse Recovery
Yi Li . . . 10

Algorithms for Massive Graphs via Linear Sketches
Andrew McGregor . . . 11

Bloom Filters in Adversarial Environments
Moni Naor . . . 11

Heavy Hitters via Cluster-Preserving Clustering
Jelani Nelson . . . 12

Approximate Near Neighbors for General Symmetric Norms
Ilya Razenshteyn . . . 12

Dynamic Space Efficient Hash Tables
Peter Sanders . . . 13

An Empirical Perspective on Locality-Sensitive Hashing
Ludwig Schmidt . . . 13

Locality Sensitive Hashing in Wild
Anshumali Shrivastava . . . 14

Locality-Sensitive Hashing of Curves
Francesco Silvestri 14

Streaming Complexity of Approximate Pattern Matching
Tatiana Starikovskaya 15

Sample(x)=(a*x<=t) is a Distinguisher with Probability 1/8
Mikkel Thorup 15

Cuckoo Hashing with Overlapping Buckets
Stefan Walzer 16

Anecdotes on Hashing
Udi Wieder 16

Efficient Algorithms for Streaming Datasets with Near-Duplicates
Qin Zhang 17

Open Problems

The Longest Chain with $(ax + b)$ mod p Hashing
Mathias Bæk Tejs Knudsen 17

Voronoi Choice Games
Michael Mitzenmacher 18

Space Complexity of Monotone Minimal Perfect Hashing
Rasmus Pagh 19

Approximate Pattern Matching in a Stream
Tatiana Starikovskaya 19

Participants 21

3 Overview of Talks

3.1 Beyond Locality Sensitive Hashing

Alexandr Andoni (Columbia University – New York, US)

Joint work of Piotr Indyk, Thijs Laarhoven, Huy Nguyen, Sasho Nikolov, Ilya Razenshteyn, Ludwig Schmidt, Negev Shekel-Nosatzki, Erik Waingarten

This talk will survey some recent work on approximate nearest neighbor search since 2014, when data-dependent hashing was introduced for the problem.

References

1 Alexandr Andoni, Thijs Laarhoven, Ilya P. Razenshteyn, Erik Waingarten: Optimal Hashing-based Time-Space Trade-offs for Approximate Near Neighbors. SODA 2017: 47-66
2 Alexandr Andoni, Ilya P. Razenshteyn, Negev Shekel Nosatzki: LSH Forest: Practical Algorithms Made Theoretical. SODA 2017: 67-78
3 Alexandr Andoni, Ilya P. Razenshteyn: Optimal Data-Dependent Hashing for Approximate Near Neighbors. STOC 2015: 793-801
4 Alexandr Andoni, Piotr Indyk, Thijs Laarhoven, Ilya P. Razenshteyn, Ludwig Schmidt: Practical and Optimal LSH for Angular Distance. NIPS 2015: 1225-1233

3.2 Distance-Sensitive Hashing

Martin Aumüller (IT University of Copenhagen, DK)

Joint work of Martin Aumüller, Tobias Christiani, Rasmus Pagh, Francesco Silvestri
Main reference M. Aumüller, T. Christiani, R. Pagh, F. Silvestri, "Distance-sensitive hashing", arXiv:1703.07867 [cs.DS], 2017.
URL http://arxiv.org/abs/1703.07867

We initiate the study of *distance-sensitive hashing*, a generalization of locality-sensitive hashing that seeks a family of hash functions such that the probability of two points having the same hash value is a given function of the distance between them. More precisely, given a distance space (X, dist) and a "collision probability function" (CPF) $f : \mathbb{R} \to [0, 1]$ we seek a distribution over pairs of functions (h, g) such that for every pair of points $x, y \in X$ the collision probability is $\Pr[h(x) = g(y)] = f(\mathrm{dist}(x, y))$. Locality-sensitive hashing is the study of how fast a CPF can *decrease* as the distance grows. For many spaces f can be made exponentially decreasing even if we restrict attention to the symmetric case where $g = h$. In this paper we study how *asymmetry* makes it possible to achieve CPFs that are, for example, increasing or unimodal. Our original motivation comes from *annulus queries* where we are interested in searching for points at distance approximately r from a query point, but we believe that distance-sensitive hashing is of interest beyond this application.

3.3 Applications of Hashing in Semantic Search

Hannah Bast (Universität Freiburg, DE)

Joint work of Hannah Bast, Florian Bäurle, Björn Buchhold, Elmar Haußmann
Main reference H. Bast, F. Bäurle, B. Buchhold, E. Haußmann, "Easy access to the freebase dataset", in Proc. of the 23rd International World Wide Web Conference (WWW 2014), pp. 95–98, ACM, 2014.
URL http://doi.acm.org/10.1145/2567948.2577016

I will briefly introduce several of our systems for semantic search on (very large) knowledge bases and text, including several live demonstrations. Each of these systems critically rely on very large hash maps, and I will explain which properties are critical for which application. For one of these requirement profiles, we will live-code a relatively simple hash map together (in C++) and show that it beats the hash map from the standard template library by a large margin.

3.4 Set Similarity Search Beyond MinHash

Tobias Christiani (IT University of Copenhagen, DK)

Joint work of Tobias Christiani, Rasmus Pagh
Main reference T. Christiani, R. Pagh, "Set similarity search beyond MinHash", in Proc. of the 49th Annual ACM SIGACT Symp. on Theory of Computing (STOC 2017), pp. 1094–1107, ACM, 2017.
URL http://doi.acm.org/10.1145/3055399.3055443

We consider the problem of approximate set similarity search under Braun-Blanquet similarity $B(x, y) = |x \cap y| / \max(|x|, |y|)$. The (b_2, b_2)-approximate Braun-Blanquet similarity search problem is to preprocess a collection of sets P such that, given a query set q, if there exists $x \in P$ with $B(q, x) \geq b_1$, then we can efficiently return $x' \in P$ with $B(q, x') > b_2$. We present a simple data structure that solves this problem with space usage $O(n^{1+\rho} \log n + \sum_{x \in P} |x|)$ and query time $O(|q| n^{\rho} \log n)$ where $n = |P|$ and $\rho = \log(1/b_1)/\log(1/b_2)$. Making use of existing lower bounds for locality-sensitive hashing by O'Donnell et al. (TOCT 2014) we show that this value of ρ is tight across the parameter space, i.e., for every choice of constants $0 < b_2 < b_1 < 1$. In the case where all sets have the same size our solution strictly improves upon the value of ρ that can be obtained through the use of state-of-the-art data-independent techniques in the Indyk-Motwani locality-sensitive hashing framework (STOC 1998) such as Broder's MinHash (CCS 1997) for Jaccard similarity and Andoni et al.'s cross-polytope LSH (NIPS 2015) for cosine similarity. Surprisingly, even though our solution is data-independent, for a large part of the parameter space we outperform the currently best data-dependent method by Andoni and Razenshteyn (STOC 2015).

3.5 BDDs for Minimal Perfect Hashing: Merging Two State-Space Compression Techniques

Stefan Edelkamp (Universität Bremen, DE)

This talk will merge two different lines of research, namely

a) state-space exploration with binary decision diagrams (BDDs), that was initially proposed for Model Checking and still is state-of-the-art in AI Planning.

b) state-space compaction with (minimal) perfect hashing, which is used in the algorithm community as a memory-based index for big data (often residing on disk).

I will show how BDDs can serve as the internal representation of a perfect hash function with linear-time ranking and unranking, and how it can be used as a static dictionary and an alternative to the recent compression schemes exploiting hypergraph theory. This will also result in a simple method to split a BDD in parts of equal number of satisfying assignments and to generate random inputs for any function represented as a BDD. As a surplus, the BDD hash function is monotone.

In terms of applications, symbolic exploration with BDD constructs a succinct representation of the state space. For each layer of the search, a BDDs is generated and stored, and will later serve as an index to do extra work like the classification of game states. Based on this approach we will show, how to strongly solve Connect-Four in a combination of symbolic and explicit-state space exploration.

3.6 Synchronization Strings: Near-Optimal Codes for Insertions and Deletions

Bernhard Haeupler (Carnegie Mellon University – Pittsburgh, US)

Joint work of Bernhard Haeupler, Amirbehshad Shahrasbi
Main reference B. Haeupler, A. Shahrasbi, "Synchronization strings: codes for insertions and deletions approaching the Singleton bound", in Proc. of the 49th Annual ACM SIGACT Symp. on Theory of Computing (STOC 2017), pp. 33–46, ACM, 2017.
URL https://doi.org/10.1145/3055399.3055498

We introduce synchronization strings as a novel way of efficiently dealing with synchronization errors, i.e., insertions and deletions. Synchronization errors are strictly more general and much harder to deal with than commonly considered half-errors, i.e., symbol corruptions and erasures. For every $\varepsilon > 0$, synchronization strings allow to index a sequence with an $\varepsilon - O(1)$ size alphabet such that one can efficiently transform k synchronization errors into $(1 + \varepsilon)k$ half-errors. This powerful new technique has many applications. In this paper, we focus on designing insdel codes, i.e., error correcting block codes (ECCs) for insertion deletion channels.

While ECCs for both half-errors and synchronization errors have been intensely studied, the later has largely resisted progress. Indeed, it took until 1999 for the first insdel codes with constant rate, constant distance, and constant alphabet size to be constructed by Schulman and Zuckerman. Insdel codes for asymptotically large or small noise rates were given in 2016 by Guruswami et al. but these codes are still polynomially far from the optimal rate-distance tradeoff. This makes the understanding of insdel codes up to this work equivalent to what was known for regular ECCs after Forney introduced concatenated codes in his doctoral thesis 50 years ago.

A direct application of our synchronization strings based indexing method gives a simple black-box construction which transforms any ECC into an equally efficient insdel code with a slightly larger alphabet size. This instantly transfers much of the highly developed understanding for regular ECCs over large constant alphabets into the realm of insdel codes. Most notably, we obtain efficient insdel codes which get arbitrarily close to the optimal rate-distance tradeoff given by the Singleton bound for the complete noise spectrum.

3.7 On-the-Fly Array Initialization in Less Space

Torben Hagerup (Universität Augsburg, DE)

Joint work of Torben Hagerup, Frank Kammer

We show that for all given $n, t, w \in \{1, 2, \dots\}$ with $n < 2^w$, an array of n entries of w bits each can be represented on a word RAM with a word length of w bits in at most $nw + \lceil n(t/(2w))^t \rceil$ bits of uninitialized memory to support constant-time initialization of the whole array and $O(t)$-time reading and writing of individual array entries. At one end of this tradeoff, we achieve initialization and access (i.e., reading and writing) in constant time with $nw + \lceil n/w^t \rceil$ bits for arbitrary fixed t, to be compared with $nw + \Theta(n)$ bits for the best previous solution, and at the opposite end, still with constant-time initialization, we support $O(\log n)$-time access with just $nw + 1$ bits, which is optimal for arbitrary access times if the initialization executes fewer than n steps.

3.8 An Adaptive Sublinear Time Block Sparse Fourier Transform

Michael Kapralov (EPFL – Lausanne, CH)

Joint work of Michael Kapralov, Volkan Cevher, Jonathan Scarlett, Amir Zandieh
Main reference V. Cevher, M. Kapralov, J. Scarlett, A. Zandieh, "An adaptive sublinear-time block sparse fourier transform", in Proc. of the 49th Annual ACM SIGACT Symp. on Theory of Computing (STOC 2017), pp. 702–715, ACM, 2017.
URL http://doi.acm.org/10.1145/3055399.3055462

The problem of approximately computing a small number k of dominant Fourier coefficients of a vector of length n quickly, and using few samples in time domain, is known as the Sparse Fourier Transform (sparse FFT) problem. A long line of work on the sparse FFT has resulted in algorithms with $O(k \log n \log(n/k))$ runtime and $O(k \log n)$ sample complexity. These results are proved using non-adaptive algorithms, and the latter sample complexity result is essentially the best possible under the sparsity assumption alone: It is known that even adaptive algorithms must use $\Omega((k \log(n/k))/ \log \log n)$ samples. By *adaptive*, we mean being able to exploit previous samples in guiding the selection of further samples.

In this work we revisit the sparse FFT problem with the added twist that the sparse coefficients approximately obey a (k_0, k_1)-block sparse model. In this model, signal frequencies are clustered in k_0 intervals with width k_1 in Fourier space, where $k = k_0 k_1$ is the total sparsity. Signals arising in applications are often well approximated by this model with $k_0 \ll k$.

We give the first sparse FFT algorithm for (k_0, k_1)-block sparse signals with the sample complexity of $O^*(k_0 k_1 + k_0 \log(1 + k_0) \log n)$ at constant signal-to-noise ratios, and sublinear runtime. A similar sample complexity was previously achieved in the works on *model-based compressive sensing* using random Gaussian measurements, but used $\Omega(n)$ runtime. To the best of our knowledge, our result is the first sublinear-time algorithm for model based compressed sensing, and the first sparse FFT result that goes below the $O(k \log n)$ sample complexity bound.

3.9 Locality Sensitive Distortion

Ravi Kumar (Google Research – Mountain View, US)

Joint work of Flavio Chierichetti, Ravi Kumar, Alessandro Panconesi, Erisa Terolli
Main reference F. Chierichetti, R. Kumar, A. Panconesi, E. Terolli, "The Distortion of Locality Sensitive Hashing", in Proc. of the 8th Innovations in Theoretical Computer Science Conference (ITCS 2017), LIPIcs, Vol. 67, pp. 54:1–54:18, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
URL https://doi.org/10.4230/LIPIcs.ITCS.2017.54

Given a pairwise similarity notion between objects, locality sensitive hashing (LSH) aims to construct a hash function family over the universe of objects such that the probability two objects hash to the same value is their similarity. LSH is a powerful algorithmic tool for large-scale applications and much work has been done to understand LSHable similarities, i.e., similarities that admit an LSH. Our work focuses on similarities that are provably non-LSHable. We propose a notion of distortion to capture the approximation of such a similarity by a similarity that is LSHable. We consider several well-known non-LSHable similarities and show tight upper and lower bounds on their distortion.

3.10 Improved ℓ_2/ℓ_2 Sparse Recovery

Yi Li (Nanyang TU – Singapore, SG)

We study the 'for-each' version of the compressed sensing and related heavy hitters problem: (1) For the compressed sensing problem, under the strongest ℓ_2/ℓ_2 error guarantee, we provide a simple scheme that uses $O(klog(n/k))$ measurements, has $O(klog^3 n)$ decoding time, and achieves $\max\{e^{-k/log^3 k}, (n/k)^{-logk}\}$ failure probability. Our result simultaneously improves the previous best scheme of Gilbert et al. (SICOMP'12) in terms of failure probability, decoding time, and column sparsity. A followup work of Gilbert at al. (ICALP'13) focuses on the very low error probability regime, and our work also improves the number of measurements and failure probability of that scheme, while achieving a similar $O(k^2 polylogn)$ decoding time. A further consequence of our arguments is that we completely resolve the measurement complexity of ℓ_2/ℓ_2 sparse recovery in the popular, average-case spiked-covariance model, providing both a new upper bound and a new lower bound. (2) For the related heavy hitters problem, under the strongest ℓ_infty/ℓ_2 error guarantee, we provide the first optimal measurement lower bound in terms of the approximation factor eps, n, and the failure probability delta, for the full range of such parameters. Our lower bound shows that the classical Count-Sketch data structure is optimal in all parameters.

3.11 Algorithms for Massive Graphs via Linear Sketches

Andrew McGregor (University of Massachusetts – Amherst, US)

URL https://people.cs.umass.edu/ mcgregor/slides/17-dagstuhl.pdf

In this talk, we will survey recent work on using random linear projections, a.k.a. sketches, to analyze massive graphs. Sketches are useful in a variety of computational models including the dynamic graph stream model were the input is defined by a stream of edge insertions and deletions that need to be processed in small space. A large number of problems have now been studied in this model including edge and vertex connectivity, spectral sparsification, triangle counting, densest subgraph, correlation clustering, vertex cover, and matching.

3.12 Bloom Filters in Adversarial Environments

Moni Naor (Weizmann Institute – Rehovot, IL)

Joint work of Moni Naor, Eylon Yogev
Main reference M. Naor, E. Yogev, "Bloom Filters in Adversarial Environments", IACR Cryptology ePrint Archive, p. 543, IACR, 2015.
URL https://eprint.iacr.org/2015/543

Many efficient data structures use randomness, allowing them to improve upon deterministic ones. Usually, their efficiency and/or correctness are analyzed using probabilistic tools under the assumption that the inputs and queries are independent of the internal randomness of the data structure. In this talk, we consider data structures in a more robust model, which we call the adversarial model. Roughly speaking, this model allows an adversary to choose inputs and queries adaptively according to previous responses. Specifically, we consider Bloom filters and prove a tight connection between Bloom filters in this model and cryptography.

A Bloom filter represents a set S of elements approximately, by using fewer bits than a precise representation. The price for succinctness is allowing some errors: for any x in S it should always answer 'Yes', and for any x not in S it should answer 'Yes' only with small probability.

In the adversarial model, we consider both efficient adversaries (that run in polynomial time) and computationally unbounded adversaries that are only bounded in the amount of queries they can make. For computationally bounded adversaries, we show that non-trivial (memory-wise) Bloom filters exist if and only if one-way functions exist. For unbounded adversaries we show that there exists a Bloom filter for sets of size n and error eps, that is secure against t queries and uses only $O(n \cdot \log(1/\varepsilon) + t)$ bits of memory. In comparison, $n \cdot \log(1/\varepsilon)$ is the best possible even under a non-adaptive adversary.

3.13 Heavy Hitters via Cluster-Preserving Clustering

Jelani Nelson (Harvard University – Cambridge, US)

Joint work of Kasper Green Larsen, Jelani Nelson, Huy L. Nguyen, Mikkel Thorup
Main reference K. Green Larsen, J. Nelson, H. L. Nguyen, M. Thorup, "Heavy Hitters via Cluster-Preserving Clustering", in Proc. of the 57th Annual Symposium on Foundations of Computer Science (FOCS 2016), pp. 61–70, IEEE, 2017.
URL https://doi.org/10.1109/FOCS.2016.16

In the "heavy hitters" or "frequent items" problem, one must process a stream of items and report those items that occur frequently. For example, a telecommunications company may wish to find popular destination IP addresses in a packet stream across one of their links, or a search engine may wish to report popular query words. A more general problem is when there are two streams and we must report those items whose frequencies significantly deviate between them; the former problem is a special case since we can artificially pretend the first of the two streams the empty stream. Such a problem naturally arises in trend detection and anomaly detection. Several algorithms were known in the literature solving these problems, such as the CountMin sketch, CountSketch, Hierarchical CountSketch and others. The goal in designing such algorithms is (1) to guarantee finding frequent items for as lax a definition of "frequent" as possible while still limiting output size, and while having low (2) memory consumption, (3) processing time per stream item, (4) query time to report the list of frequent times, and (5) failure probability. Previous solutions could perform well on various subsets of these metrics, but not on all 5 simultaneously.

We design a new algorithm, ExpanderSketch, which performs well on all 5. Our main innovation is a novel reduction to a new graph-clustering problem we formulate, in which finding most of every cluster guarantees finding all the frequent items. We then solve this problem by devising a novel spectral clustering algorithm potentially of independent interest, based on divide-and-conquer and local search.

3.14 Approximate Near Neighbors for General Symmetric Norms

Ilya Razenshteyn (MIT – Cambridge, US)

Joint work of Alexandr Andoni, Huy L. Nguyen, Aleksandar Nikolov, Ilya Razenshteyn, Erik Waingarten
Main reference A. Andoni, H. L. Nguyen, A. Nikolov, I. P. Razenshteyn, E. Waingarten, "Approximate near neighbors for general symmetric norms", in Proc. of the 49th Annual ACM SIGACT Symp. on Theory of Computing (STOC 2017), pp. 902–913, ACM, 2017.
URL http://dx.doi.org/10.1145/3055399.3055418

We show that every *symmetric* normed space admits an efficient nearest neighbor search data structure with *doubly-logarithmic* approximation. Specifically, for every n, $d = n^{o(1)}$, and every d-dimensional symmetric norm $\|\cdot\|$, there exists a data structure for $\mathrm{poly}(\log\log n)$-approximate nearest neighbor search over $\|\cdot\|$ for n-point datasets achieving $n^{o(1)}$ query time and $n^{1+o(1)}$ space. The main technical ingredient of the algorithm is a low-distortion embedding of a symmetric norm into a low-dimensional iterated product of top-k norms.

We also show that our techniques cannot be extended to *general* norms.

3.15 Dynamic Space Efficient Hash Tables

Peter Sanders (KIT – Karlsruher Institut für Technologie, DE)

Joint work of Tobias Maier, Peter Sanders
Main reference T. Maier, P. Sanders, "Dynamic Space Efficient Hashing", in Proc. of the 25th Annual European Symposium on Algorithms (ESA 2017), LIPIcs, Vol. 87, pp. 58:1–58:14, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
URL https://doi.org/10.4230/LIPIcs.ESA.2017.58

We consider space efficient hash tables that can grow and shrink dynamically and are always highly space efficient, i.e., their space consumption is always close to the lower bound even while growing and when taking into account storage that is only needed temporarily. None of the traditionally used hash tables have this property. We show how known approaches like linear probing and bucket cuckoo hashing can be adapted to this scenario by subdividing them into many subtables or using virtual memory overcommitting. However, these rather straightforward solutions suffer from slow amortized insertion times due to frequent reallocation in small increments.

Our main result is DySECT (**Dy**namic **S**pace **E**fficient **C**uckoo **T**able) which avoids these problems. DySECT consists of many subtables which grow by doubling their size. The resulting inhomogeneity in subtable sizes is equalized by the flexibility available in bucket cuckoo hashing where each element can go to several buckets each of which containing several cells. Experiments indicate that DySECT works well with load factors up to 98%. With up to 2.7 times better performance than the next best solution.

3.16 An Empirical Perspective on Locality-Sensitive Hashing

Ludwig Schmidt (MIT – Cambridge, US)

Joint work of Alexandr Andoni, Piotr Indyk, Thijs Laarhoven, Ilya Razenshteyn, Ludwig Schmidt, Kunal Talwar
Main reference A. Andoni, P. Indyk, T. Laarhoven, I. Razenshteyn, L. Schmidt, "Practical and Optimal LSH for Angular Distance", in Proc. of the 28th Annual Conference on Advances in Neural Information Processing Systems (NIPS 2015), pp. 1225–1233, 2015.
URL http://papers.nips.cc/paper/5893-practical-and-optimal-lsh-for-angular-distance

Locality-Sensitive Hashing offers attractive theoretical guarantees for (approximate) nearest neighbor search. In this talk, I review recent empirical work on nearest neighbor algorithms and how it compares with LSH-based methods. I then describe our work on the cross-polytope hash, which combines good theoretical properties with good practical performance. Finally, I show some experiments with locality-sensitive hashing in the context of deep neural networks.

3.17 Locality Sensitive Hashing in Wild

Anshumali Shrivastava (Rice University – Houston, US)

I will talk about some of my wild explorations with probabilistic hashing algorithms and some of my very recent findings. It turns out there is another feather in the cap for LSH. It can be used for smart sampling and estimations. If we view (K, L) LSH algorithm that adaptively samples (very efficiently) data x with probability $1 - (1 - p^K)^L$ then this can be turned into efficient unbiased estimations algorithms. (p being the collision probability between query and data x). The similar sampling idea can be used to significantly reduce the computations in classical machine learning algorithms such Deep Learning (using our recent success with asymmetric hashing for inner products). In another wild exploration, I will show how Minhash is a better hashing scheme for cosine similarity than Simhash. I will highlight the computational bottleneck, i.e. the hashing time, and will show an efficient variant of minwise hashing. If time permits, I will demonstrate the use of probabilistic hashing for obtaining practical privacy-preserving algorithms.

3.18 Locality-Sensitive Hashing of Curves

Francesco Silvestri (University of Padova, IT)

Joint work of Anne Driemel, Francesco Silvestri
Main reference A. Driemel, F. Silvestri, "Locality-Sensitive Hashing of Curves", in Proc. of the 33rd Int'l Symposium on Computational Geometry (SoCG 2017), LIPIcs, Vol. 77, pp. 37:1-37:16, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
URL https://doi.org/10.4230/LIPIcs.SoCG.2017.37

We study data structures for storing a set of polygonal curves in $\Re^d$ such that, given a query curve, we can efficiently retrieve similar curves from the set, where similarity is measured using the discrete Frechet distance or the dynamic time warping distance. To this end we devise the first locality-sensitive hashing schemes for these distance measures. A major challenge is posed by the fact that these distance measures internally optimize the alignment between the curves. We give solutions for different types of alignments including constrained and unconstrained versions. For unconstrained alignments, we improve over a result by Indyk from 2002 for short curves. Let n be the number of input curves and let m be the maximum complexity of a curve in the input. In the particular case where $m \leq \frac{\alpha}{4d} \log n$, for some fixed $\alpha > 0$, our solutions imply an approximate near-neighbor data structure for the discrete Frechet distance that uses space in $O(n^{1+\alpha} \log n)$ and achieves query time in $O(n^\alpha \log^2 n)$ and constant approximation factor. Furthermore, our solutions provide a trade-off between approximation quality and computational performance: for any parameter $k \in [m]$, we can give a data structure that uses space in $O(2^{2k} m^{k-1} n \log n + nm)$, answers queries in $O(2^{2k} m^k \log n)$ time and achieves approximation factor in $O(m/k)$.

3.19 Streaming Complexity of Approximate Pattern Matching

Tatiana Starikovskaya (University Paris-Diderot, FR)

Joint work of Raphaël Clifford, Tatiana Starikovskaya
Main reference R. Clifford, T. Starikovskaya, "Approximate Hamming distance in a stream", in Proc. of the 43rd Int'l Colloquium on Automata, Languages, and Programming (ICALP 2016), LIPIcs, Vol. 55, pp. 20:1–20:14, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2016.
URL http://dx.doi.org/10.4230/LIPIcs.ICALP.2016.20
Main reference T. Starikovskaya, "Streaming and communication complexity of approximate pattern matching", in Proc. of the 28th Symposium on Combinatorial Pattern Matching (CPM 2017), LIPIcs, Vol. 78, pp. 13:1–13:11, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
URL http://dx.doi.org/10.4230/LIPIcs.CPM.2017.13

In the approximate pattern matching problem we are given two strings, a pattern and a text, and must find all approximate occurrences of the pattern in the text. An approximate occurrence of the pattern is a factor of the text such that the distance (Hamming, edit, ...) between it and the pattern is small. This problem is a fundamental problem of text processing and has myriad of applications. However, classical solutions to this problem cannot be used for processing massive data as they demonstrate unfavourable space lower bounds. In the talk I will present several new algorithms for this problem under the streaming model of computation. In this model the input is received as a stream, one item at a time, and the algorithms are not allowed to store a copy of the input without accounting for it, which leads to particularly space-efficient solutions.

3.20 Sample(x)=(a*x<=t) is a Distinguisher with Probability 1/8

Mikkel Thorup (University of Copenhagen, DK)

Main reference M. Thorup, "Sample (x) = (a*x<=t) is a Distinguisher with Probability 1/8", in Proc. of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS 2015), pp. 1277–1291, IEEE, 2015.
URL http://dx.doi.org/10.1109/FOCS.2015.82

A random sampling function Sample : $U \to \{0, 1\}$ for a key universe U is a distinguisher with probability P if for any given assignment of values $v(x)$ to the keys $x \in U$, including at least one non-zero $v(x) \neq 0$, the sampled sum, that is, $\text{sum}\{v(x)|x \in U, \text{Sample}(x) = 1\}$, is non-zero with probability at least P. Here the key values may come from any commutative monoid (addition is commutative and associative and zero is neutral). Such distinguishers were introduced by Vazirani [PhD thesis 1986], and Naor and Naor used them for their small bias probability spaces [STOC'90]. Constant probability distinguishers are used for testing in contexts where the key values are not computed directly, yet where the sum is easily computed. A simple example is when we get a stream of key value pairs $(x_1, v_1), (x_2, v_2), ..., (x_n, v_n)$ where the same key may appear many times. The accumulated value of key x is $v(x) = \text{sum}\{v_i \mid x_i = x\}$. For space reasons, we may not be able to maintain $v(x)$ for every key x, but the sampled sum is easily maintained as the single value $\text{sum}\{v_i \mid \text{Sample}(x_i)\}$.

Here we show that when dealing with w-bit integers, if a is a uniform odd w-bit integer and t is a uniform w-bit integer, then $\text{Sample}(x) = [ax \mod 2^w \leq t]$ is a distinguisher with probability 1/8. Working with standard units, that is, w=8, 16, 32, 64, we exploit that w-bit multiplication works modulo 2^w, discarding overflow automatically, and then the sampling decision is implemented by the C-code `a * x <= t`. Previous such samplers were much less computer friendly, e.g., the distinguisher of Naor and Naor [STOC'90] was more complicated and involved a 7-independent hash function.

3.21 Cuckoo Hashing with Overlapping Buckets

Stefan Walzer (TU Ilmenau, DE)

Cuckoo hashing is the task of finding an injective assignment of m objects to n memory positions, where each object e is only allowed to reside in one of $k \geq 2$ different positions that are chosen for e independently and uniformly at random. It is well understood how the number of choices k affects the critical density $c^* \in (0, 1)$ around which the probability that a valid assignment exists jumps from almost 0 to almost 1.

In order to profit from memory caches when searching for an object, it is reasonable in practice to weaken the requirement that all choices be independent, and instead assign each object k contiguous memory regions, each of size ℓ (totalling $k \cdot \ell$ choices) that can be searched with k cache faults only. In previous work, the possible memory regions formed a partition of the total memory available. Dietzfelbinger and Weidling noticed that space utilisation can be improved if memory regions are allowed to overlap imperfectly, i.e. the choices for each object are the union of k intervals of length ℓ, where the start of an interval is not necessarily a multiple of ℓ. We confirm the experimental results by deriving corresponding thresholds rigorously.

Our main tool is a theorem due to Lelarge, Leconte and Massoulié which uses methods from statistical physics. It allows to derive the thresholds from properties of infinite trees that the underlying hypergraphs converge to.

3.22 Anecdotes on Hashing

Udi Wieder (VMware – Palo Alto, US)

In this talk I will present some anecdotes encountered and some lessons learned from attempts to implement hash tables and approximate set representations in industry. My experience involves implementations with high level languages and not extremely large data sets, and as such differs from the lessons learned in big-data companies.

3.23 Efficient Algorithms for Streaming Datasets with Near-Duplicates

Qin Zhang (Indiana University – Bloomington, US)

Joint work of Djamal Belazzougui, Di Chen, Jiecao Chen, Haoyu Zhang, Qin Zhang

In this talk I will discuss how to analyze data streams with near-duplicates. I will talk about two specific problems: (1) the estimation of the number of distinct elements, and (2) similarity join under edit distance. The messages I would like to deliver are: (1) there is an algorithmic framework for estimating the number of distinct elements in the streaming model, given that the underlying metric space admit a locality sensitive hashing (LSH) with a special property. (2) For those metric spaces X that do not admit LSH (e.g., the edit distance), we can first try to embed X to another metric space which has good LSHs. (3) If we want to handle exact distance thresholds (for defining near-duplicates) then we may need to use sketches.

References

1 Di Chen and Qin Zhang, "Streaming Algorithms for Robust Distinct Elements", SIGMOD 2016
2 Djamal Belazzougui and Qin Zhang, "Edit Distance: Sketching, Streaming and Document Exchange", FOCS 2016
3 Haoyu Zhang and Qin Zhang, "EmbedJoin: Efficient Edit Similarity Joins via Embeddings", KDD 2017

4 Open Problems

4.1 The Longest Chain with $(ax + b) \bmod p$ Hashing

Mathias Bæk Tejs Knudsen (University of Copenhagen, DK)

Let p, n be integers with $p > n$ and p a prime. Let $h : \{0, 1, \ldots, p-1\} \to \{0, 1, \ldots, n-1\}$ be the hash function defined by $h(x) = ((ax + b) \bmod p) \bmod n$ where $a, b \in \{0, 1, \ldots, p-1\}$ are chosen independently and uniformly random. Consider using h to build a hash table of size n with n keys. That is, we let $X \subset \{0, 1, \ldots, p-1\}$ be a set of size n, and for $x \in X$ we put x into the i'th linked list in the table if $h(x) = i \in \{0, 1, \ldots, n-1\}$. We let $L(X)$ denote the length of the longest chain and note $L(X)$ is a random variable defined by:

$$L(X) = \max_{i \in \{0,1,\ldots,n-1\}} \{x \in X \mid h(x) = i\} .$$

The worst case time it takes to search for an element in the hash table is $\Theta(L(X))$, which motivates the study of $\mathrm{E}[L(X)]$. We let M denote the maximal expected value over all choices of X, i.e.

$$M = \max \{\mathrm{E}[L(X)] \mid X \subset \{0, 1, \ldots, p-1\}, |X| = n\} .$$

We are interested in finding the value of M.

Open question: What is the smallest non-negative real number c such that $M \leq n^{c+o(1)}$. Is it true that $c > 0$? Is it true that $c < \frac{1}{3}$? (It is an easy exercise to prove that such a c exists)

Known results: It is known that $c \in \left[0, \frac{1}{3}\right]$. The upper bound is from [1].

References

1 Knudsen, M. B. T. Linear Hashing is Awesome. Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS), pp. 345-352. 2016.

4.2 Voronoi Choice Games

Michael Mitzenmacher (Harvard University – Cambridge, US)

Consider the unit torus. We consider a game with n players (we think of n as growing large). Each player is given $m > 1$ points on the torus chosen independently and uniformly at random, giving mn total points. (We think of m as a small constant; for example, consider the case $m = 2$.) You may assume all players know all of points given to all of the players. In the game, each player must simultaneously choose one of their m points; using those n points, we construct a Voronoi diagram, and each player obtains a score corresponding to the size of the area of the cell of the Voronoi diagram containing their point.

Open question: Is there with probability $1 - o(1)$ a pure Nash equilibrium (e.g., a selection of points for all of the players where no single individual player wants to change their decision to another point) when the goal is for each player to maximize (respectively minimize) their score?

Known results: There is little or nothing known about this specific variant, but one can come up with all sorts of variations of the problem; see [1] for details, and in particular for how individual players having their own sets of points differentiates these problems from previous Voronoi games. A more readily analyzed variant – but still very difficult – has each player obtaining points chosen uniformly at random on the boundary of the unit circle, and after each player selects one of their points, their score corresponds to the length of the clockwise arc starting at their point until the next point for another player is reached. It is known that the number of pure Nash equilibrium when each player tries to maximize their arc length in this setting is between $(0.19)^{m-1}m$ and m; for even this simpler variant, these results do not settle the question of whether a pure Nash equilibrium exists with high probability.

References

1 Boppana, M., Hod, R., Mitzenmacher, M., Morgan, T. Voronoi Choice Games. ICALP 2016: 23:1-23:13.

4.3 Space Complexity of Monotone Minimal Perfect Hashing

Rasmus Pagh (IT University of Copenhagen, DK)

Given a set S of n elements from $\{1, \ldots, u\}$ we wish to store a function f such that for each $x \in S$, $f(x) = \text{rank}_S(x)$, where $\text{rank}_S(x) = |\{y \in S \mid y \leq x\}|$. Such a function is referred to as a monotone minimal perfect hash function (MMPHF) for S. We want a family F of smallest possible size such that for every S of size n there exists $f \in F$ that is a MMPHF for S. This will allow us to store f in $\log_2 |F|$ bits of space. To simplify the problem, you may assume free access to a random function $h\colon \mathbb{N} \to \{0, 1\}$; the space usage should then hold with high probability over the randomness in h.

Open question: What is the best possible value of $\log_2 |F|$ for an MMPHF family F?

Known results: Lower bound of $\Omega(n)$ bits [2]. Holds even for minimal perfect hashing without monotonicity. Upper bound of $O(n \min(\log n, \log\log\log u))$ bits [1].

References

1 Belazzougui, D., Boldi, P., Pagh, R., & Vigna, S. Monotone minimal perfect hashing: searching a sorted table with $O(1)$ accesses. In Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2009), pages 785-794, 2009. SIAM.

2 Mehlhorn, K. On the program size of perfect and universal hash functions. In Proceedings of the 23th Annual Symposium on Foundations of Computer Science (FOCS 1982), pages 170–175, 1982. IEEE.

4.4 Approximate Pattern Matching in a Stream

Tatiana Starikovskaya (University Paris-Diderot, FR)

In the approximate pattern matching problem we are given a text T and a pattern P of length n, and we must decide for each prefix of T if it ends with a string that is at the edit distance at most k from P. If this is the case, we must output the edit distance and the corresponding edit operations.

Open question: What is the complexity of the problem in the streaming model? (In the streaming model we assume that the pattern and the text arrive as streams, first the pattern, then the stream. The space is all the space used by the algorithm, included the space needed to store the information about the input. The time is the worst-case time the algorithm spends for processing a prefix of the text.)

Known results: There is a streaming algorithm that uses $O(k^8\sqrt{n} log^6 n)$ space and $O((k^2\sqrt{n} + k^{13}) \log n)$ time per arrival[1]. It is the first sublinear-space algorithm for this well-known problem. The main tool is a new sketch for computing the edit distance between two strings suggested by Belazzougui and Zhang in FOCS 2016 [2].

References

1 Tatiana Starikovskaya. Communication and streaming complexity of approximate pattern matching. To appear in CPM 2017.

2 Djamal Belazzougui, Qin Zhang: Edit Distance: Sketching, Streaming, and Document Exchange. FOCS 2016: 51-60.

Participants

- Alexandr Andoni
Columbia University – New York, US
- Martin Aumüller
IT University of Copenhagen, DK
- Hannah Bast
Universität Freiburg, DE
- Djamal Belazzougui
CERIST – Algiers, DZ
- Vladimir Braverman
Johns Hopkins University – Baltimore, US
- Andrei Z. Broder
Google Inc. – Mountain View, US
- Tobias Christiani
IT University of Copenhagen, DK
- Artur Czumaj
University of Warwick – Coventry, GB
- Søren Dahlgaard
University of Copenhagen, DK
- Martin Dietzfelbinger
TU Ilmenau, DE
- Stefan Edelkamp
Universität Bremen, DE
- Tom Friedetzky
Durham University, GB
- Andreas Goerdt
TU Chemnitz, DE
- Bernhard Haeupler
Carnegie Mellon University – Pittsburgh, US
- Torben Hagerup
Universität Augsburg, DE
- Michael Kapralov
EPFL – Lausanne, CH
- Mathias Bæk Tejs Knudsen
University of Copenhagen, DK
- Ravi Kumar
Google Research – Mountain View, US
- Yi Li
Nanyang TU – Singapore, SG
- Andrew McGregor
University of Massachusetts – Amherst, US
- Michael Mitzenmacher
Harvard University – Cambridge, US
- Moni Naor
Weizmann Institute – Rehovot, IL
- Jelani Nelson
Harvard University – Cambridge, US
- Rasmus Pagh
IT University of Copenhagen, DK
- Konstantinos Panagiotou
LMU München, DE
- Ely Porat
Bar-Ilan University – Ramat Gan, IL
- Ilya Razenshteyn
MIT – Cambridge, US
- Peter Sanders
KIT – Karlsruher Institut für Technologie, DE
- Tamás Sarlós
Google Inc. – Mountain View, US
- Ludwig Schmidt
MIT – Cambridge, US
- Anshumali Shrivastava
Rice University – Houston, US
- Francesco Silvestri
University of Padova, IT
- Christian Sohler
TU Dortmund, DE
- Tatiana Starikovskaya
University Paris-Diderot, FR
- Mikkel Thorup
University of Copenhagen, DK
- Stefan Walzer
TU Ilmenau, DE
- Udi Wieder
VMware – Palo Alto, US
- Philipp Woelfel
University of Calgary, CA
- David P. Woodruff
IBM Almaden Center – San Jose, US
- Qin Zhang
Indiana University – Bloomington, US

Report from Dagstuhl Seminar 17191

Theory of Randomized Optimization Heuristics

Edited by

Carola Doerr[1], Christian Igel[2], Lothar Thiele[3], and Xin Yao[4]

1 **CNRS & UPMC, Paris, FR,** carola.doerr@mpi-inf.mpg.de
2 **University of Copenhagen, DK,** igel@diku.dk
3 **ETH Zürich, CH,** thiele@ethz.ch
4 **University of Birmingham, GB,** x.yao@cs.bham.ac.uk

Abstract

This report summarizes the talks, breakout sessions, and discussions at the Dagstuhl Seminar 17191 on *Theory of Randomized Optimization Heuristics*, held during the week from May 08 until May 12, 2017, in Schloss Dagstuhl – Leibniz Center for Informatics. The meeting is the successor of the "Theory of Evolutionary Algorithm" seminar series, where the change in the title reflects the development of the research field toward a broader range of heuristics. The seminar has hosted 40 researchers from 15 countries. Topics that have been intensively discussed at the seminar include population-based heuristics, constrained optimization, non-static parameter choices as well as connections to research in machine learning.

Seminar May 7–12, 2017 – http://www.dagstuhl.de/17191
1998 ACM Subject Classification F.2
Keywords and phrases algorithms and complexity, evolutionary algorithms, machine learning, optimization, soft computing
Digital Object Identifier 10.4230/DagRep.7.5.22
Edited in cooperation with Martin S. Krejca

1 Executive Summary

Carola Doerr
Christian Igel
Lothar Thiele
Xin Yao

Randomized search and optimization heuristics such as evolutionary algorithms, ant colony optimization, particle swarm optimization, and simulated annealing, have become established problem solvers. They have successfully been applied to a wide range of real-world applications, and they are applicable to problems that are non-continuous, multi-modal, and/or noisy as well as to multi-objective and dynamic optimization tasks. Theory of randomized optimization heuristics aims at providing mathematically founded insights into the working principles of these general-purpose problem solvers, and at developing new and more powerful heuristic optimization methods in a principled way. The seminar has covered several important streams in this research discipline. Among several other topics, extended discussions have been held on the advantages of population-based heuristics and of non-static parameter choices, optimization problems with constrains, as well as existing and possible connections to research in machine learning.

The seminar continues to be one of the key stimulator for novel ideas, tools, and approaches in the theory of randomized optimization heuristics. Accordingly, the acceptance rate for the invitations has been staying at a very high level.

Topics

The research in theory of randomized optimization heuristics is as broad as the applicability of these methods. The seminar succeeded in covering the various theoretical approaches. There was a focus on important cross-cutting topics, which we briefly outline in the following.

One of the most prominent research areas in the theory of randomized optimization heuristics deals with *runtime* and *convergence analysis*, aiming at proving bounds on the speed of the convergence to an optimal solution. Typical questions concern the advantages of certain algorithmic choices, such as

- the size of the memory (*population*),
- the usage of different sampling strategies (*variation* of previously sampled search points, in particular via *mutation* of one previously evaluated solution candidate and *recombination* of two or more previous search points), and
- the selection strategies (e.g., *elitist selection* which never discards a best-so-far solution vs. the non-elitist Boltzmann strategies found in Simulated Annealing, SSWM, and the Metropolis algorithm).

One of the most relevant objectives in empirical and theoretical works on randomized optimization heuristics is to determine the best parameter settings for the above-described algorithmic components. Given the complex interactions between the parameter values, this *parameter tuning* task is a highly challenging one. It is further complicated by the observation that for most problems the optimal parameter settings change during the optimization process, thus asking for *parameter control* mechanisms that adjust the parameter value to the current state of the optimization. Identifying such reasonable (and possibly provably optimal) ways to update the parameter choices has been one of the intensively discussed topics of the seminar. Significant progress towards a better understanding of different parameter update schemes has been obtained in the last few years, as has been demonstrated by several talks, for example on self-adaptive and self-adjusting parameter updates as well as on estimation of distribution algorithms. Among other results, several connections to related questions in machine learning have been made, motivating the organizers to include machine learning as a focus topic of this seminar.

Randomized search heuristics are currently very popular in general machine learning[1] in the form of *Bayesian optimization*. However, there has been little connection between the research in Bayesian optimization and the established work on randomized search heuristics, and the seminar was a step to change this. The first talk of the seminar was an extended introduction to Bayesian optimization by Matthew W. Hoffman from Google DeepMind, a leading expert in the field. The talk set the stage for informed discussions on similarities and differences between methods—and potential synergies between the research fields. Thompson sampling, an important algorithm in Bayesian optimization, was revisited in the talk by Jonathan Shapiro on dueling bandit problems, which demonstrated randomized search heuristics in a scenario of high commercial relevance. A common application of randomized

[1] One may well argue that randomized search heuristics actually belong to the broader field of machine learning methods.

search heuristics in general machine learning is model selection, for example finding a tailored structure of a neural network. This was addressed in the talk by Olivier Teytaud from Google Brain, who discussed model selection heuristics for large-scale machine learning systems. Randomized search heuristics are also successfully used for reinforcement learning. Arina Buzdalova presented work in which the connection is the other way round: ideas from reinforcement learning are used to improve randomized optimization (by controlling the choice of objectives).

Another intensively discussed topic, highly relevant in both discrete and continuous optimization, was constrained optimization. Here the main research questions concern the different ways to model constrained problems in black-box settings, and suitable algorithmic approaches. In addition to a number of theoretical results on constrained optimization, the need for a well-designed benchmark suite has also been discussed. As a result of one of the breakout sessions of the previous Dagstuhl Seminar 15211 on *Theory of Evolutionary Computation*, Dimo Brockhoff presented the recent extension of the COCO benchmark set (http://coco.gforge.inria.fr/doku.php) to constrained optimization. Dirk Arnold presented some work indicating that this extension of COCO is very timely, and much needed in the randomized search heuristics community. Furthermore, another breakout session has been held this year on the topic of constrained optimization, organized by Frank Neumann, with a focus on the different ways to model soft and hard constraints in discrete black-box optimization.

Organization

The seminar schedule has offered a good flexibility for the participants to propose talks and discussions of different lengths. 29 talks of 10–30 minutes each have been held in total, in the plenary room. These plenary talks were complemented by a introductory tutorial on Bayesian Optimization by Matt Hoffman on Monday morning and by 7 breakout sessions on various topics, including methodology-oriented discussions on the applicability of drift analysis in continuous domains or how to interpret the CMA-ES in the framework of information geometry optimization as well as problem-driven brainstorming on constrained optimization, the role of diversity in heuristic search, preference-based selection, and the method of estimation of distribution algorithms. Another breakout session was devoted to discussing the importance and possible obstacles of bringing theory-and practice-driven research in heuristic optimization closer together. The breakout sessions have been held on Tuesday, Wednesday, and Thursday afternoon, respectively, and have all witnessed high attendance rates. All talks and breakout sessions are summarized in Sections 3 and 4 of the present report.

We would like to express our gratitude to the Dagstuhl staff and all participants for making this Dagstuhl Seminar 17191 on *Theory of Randomized Optimization Heuristics* such a successful event, which has been a pleasure to organize.

Carola Doerr (CNRS and Pierre et Marie Curie University Paris 6, FR)
Christian Igel (University of Copenhagen, DK)
Lothar Thiele (ETH Zürich, CH)
Xin Yao (University of Birmingham, GB and SUSTech Shenzhen, CH)

2 Table of Contents

Executive Summary
Carola Doerr, Christian Igel, Lothar Thiele, and Xin Yao 22

Overview of Talks

Optimal Step-Size for the Weighted Recombination Evolution Strategy
Youhei Akimoto 28

Evolutionary Computation with Constraints
Dirk V. Arnold 28

Connecting Stability of Markov Chains and Deterministic Control Models for Analyzing Randomized Algorithms
Anne Auger 28

Towards a Theory of CMA-ES: But First, Simplify Your CMA-ES!
Hans-Georg Beyer 29

Progress Report: Towards a Constrained Test Suite for COCO
Dimo Brockhoff 29

How to Exploit Your Fitness-Distance Correlation: Runtime Analysis of the $(1 + (\lambda, \lambda))$ GA on Random Satisfiable 3-CNF Formulas
Maxim Buzdalov and Benjamin Doerr 30

Is it necessary to perform multi-objective optimization when doing multiobjectivization?
Arina Buzdalova 30

On the Variable Interaction Graph in Gray-Box Optimization
Francisco Chicano 31

Fast Genetic Algorithms
Benjamin Doerr 31

Optimal Recombination for the Asymmetric TSP: Theory and Experiment
Anton V. Eremeev 32

Mathematical Models of Artificial Genetic Representations with Neutrality
Carlos M. Fonseca 32

Monotone Functions on Bitstrings – Some Structural Notes
Christian Gießen 33

Global Convergence of the (1+1)-ES
Tobias Glasmachers 34

How to Guaranty Positive Definiteness in Active CMA-ES
Nikolaus Hansen 35

An overview of Bayesian Optimization
Matthew W. Hoffman 35

Estimation of Distribution Algorithms
Martin S. Krejca 35

Convergence in Genetic Programming
William B. Langdon 36

Landscape of the Triangle Program
William B. Langdon 36

Runtime Analysis Evolutionary Algorithms with Self-adaptive Mutation Rates
Per Kristian Lehre 37

Noise models for comparison-based evolutionary algorithms
Johannes Lengler 37

Features, Diversity, Random Walks and Digital Art
Frank Neumann 38

Standard Steady State Genetic Algorithms can Hillclimb Faster than Mutation-only Evolutionary Algorithms
Pietro S. Oliveto 39

Linear multiobjective drift analysis
Jonathan E. Rowe 39

Theoretical Aspects of the Averaged Hausdorff Indicator in Biobjective Optimization
Günter Rudolph 40

Max-Min Thompson Sampling for the K-Armed Dueling Bandit Problem
Jonathan L. Shapiro and Joseph Mellor 40

Fundamentals of ESs' Statistical Learning
Ofer M. Shir 41

Low discrepancy for one-shot optimization
Olivier Teytaud 41

Recent Advances in Runtime Analysis of Estimation-of-Distribution Algorithms
Carsten Witt 42

(1+λ) Evolutionary Algorithm with Self-Adjusting Mutation Rate
Jing Yang 43

Working groups

Breakout Session: Information Geometric Optimization
Youhei Akimoto 43

Breakout Session: Preference-based Selection in Evolutionary Multiobjective Optimisation
Carlos M. Fonseca 44

Breakout Session: Drift Theorems for Continuous Optimization
Tobias Glasmachers 45

Breakout Session: Discrete Estimation of Distribution Algorithms
Martin S. Krejca 46

Breakout Session: Theory of Evolutionary Algorithms for Problems With (Dynamic) Constraints
Frank Neumann 47

Breakout Session: Diversity
Lothar Thiele 48

Breakout Session: COST Action CA15140
Christine Zarges 49

Seminar Schedule 52

Participants 55

3 Overview of Talks

3.1 Optimal Step-Size for the Weighted Recombination Evolution Strategy

Youhei Akimoto (Shinshu University – Nagano, JP)

Joint work of Youhei Akimoto, Anne Auger, Nikolaus Hansen

We focus on the optimality of the step-size for the weighted recombination evolution strategy. In the previous work by D. Arnold, the optimal step-size and the optimal weights are derived on the Sphere function in the limit of the dimension to infinity. In this talk, we address the optimal step-size for a finite dimension, especially when the population size is close to or greater than the dimension. We derive the optimal step-size under the limit of the learning rate for the mean vector to infinity. We show that the derived optimal step-size provides a good approximation on a finite dimensional Sphere function for the standard setup of the learning rate of the mean vector update, i.e., the learning rate equal one.

3.2 Evolutionary Computation with Constraints

Dirk V. Arnold (Dalhousie University – Halifax, CA)

I discuss several aspects of constrained black-box optimization, including a taxonomy of constraints in simulation based optimization recently proposed by Le Digabel and Wild, work on attempting to arrive at an analytically based understanding of the behaviour of evolution strategies on simple constrained problems, and a comparison of algorithms on a set of test problems.

3.3 Connecting Stability of Markov Chains and Deterministic Control Models for Analyzing Randomized Algorithms

Anne Auger (INRIA RandOpf Team, FR)

Joint work of Alexandre Chotard, Anne Auger
Main reference A. Chotard, A. Auger, “Verifiable Conditions for the Irreducibility and Aperiodicity of Markov Chains by Analyzing Underlying Deterministic Models”, arXiv:1508.01644v4 [math.PR], 2017.
URL https://arxiv.org/abs/1508.01644v4

Motivated by the analysis of randomized search optimization algorithms, this talk presents connections between Markov chain theory and stability of underlying deterministic control models. We consider a general model of Markov chain $\Phi_{t+1} = F(\Phi_t, \alpha(\Phi_t, U_{t+1}))$ where $\{U_t, t > 0\}$ are i.i.d. random vectors, F is typically C^1 and $\alpha(x, U_1)$ admits a lower semi continuous density. This model embeds Markov chains that arise in stochastic optimization where α models the selection which is discontinuous but where we can derive a lower semicontinuous density.

We show that there is equivalence between the existence of a globally attracting (GA) state for the deterministic control model and the φ-irreducibility of the Markov chain provided a controllability condition holds. We then show that the support of the irreducibility measure is the set of GA states. Last we show a practical condition for proving φ-irreducibility that consists in showing the existence of a GA state where the controllability condition is satisfied.

3.4 Towards a Theory of CMA-ES: But First, Simplify Your CMA-ES!

Hans-Georg Beyer (Fachhochschule Vorarlberg – Dornbirn, AT)

Joint work of Hans-Georg Beyer, Bernhard Sendhoff
Main reference H.-G. Beyer, B. Sendhoff, “Simplify Your Covariance Matrix Adaptation Evolution Strategy”, IEEE Transactions on Evolutionary Computation, Vol. 21(5), pp. 746–759, IEEE, 2017.
URL http://dx.doi.org/10.1109/TEVC.2017.2680320

Before starting the endeavor of a theoretical convergence analysis of the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) it seems advisable to simplify this algorithm in such a manner that on the one hand it gets more amenable to such an analysis, but on the other hand without sacrificing the good optimization performance of the original CMA-ES.

In this talk it will be shown that one can remove one of the evolution path calculations from the CMA-ES keeping only the one used for the mutation strength control. In a second step it will be shown that one can also get rid of the covariance matrix update, thus removing the “C” from the CMA-ES resulting in the novel MA-ES that performs nearly equally well as the original strategy.

Besides the increased simplicity of the novel MA-ES, it also has a reduced algorithmic complexity of $\mathrm{O}(N^2)$ compared to the original $\mathrm{O}(N^3)$. Furthermore, the new M-matrix update rule derived has a special structure that allows for a direct interpretation of the M-matrix update. This update is driven by the departure of the actually selected (isotropically generated) z-vectors from isotropy. That is, the M-matrix changes its “shape” until the composition of the objective function $f(y + \sigma \cdot M \cdot z) =: g(z)$ has transformed the original problem $f(x)$ into a sphere function $g(z)$.

3.5 Progress Report: Towards a Constrained Test Suite for COCO

Dimo Brockhoff (INRIA Saclay – Île-de-France, FR)

Joint work of Anne Auger, Nikolaus Hansen, Asma Atamna, Olaf Mersmann, Tea Tusar, Dejan Tusar, Phillipe Sampaio

In the previous edition of this workshop series on Theory of Randomized Optimization Heuristics, a breakout session on “Constrained Blackbox Optimization Benchmarking” was held to discuss the first steps towards a constrained test suite within the well-known Comparing Continuous Optimizers platform (COCO, github.com/numbbo/coco/) and to identify (theoretical) questions related to this extension of the platform. In this talk, I quickly reminded us on how we benchmark optimization algorithms in COCO on unconstrained problems and then reported on the progress we made since the last breakout session towards a new constrained test suite in COCO.

3.6 How to Exploit Your Fitness-Distance Correlation: Runtime Analysis of the $(1 + (\lambda, \lambda))$ GA on Random Satisfiable 3-CNF Formulas

Maxim Buzdalov (ITMO University – St. Petersburg, RU) and Benjamin Doerr (Ecole Polytechnique – Palaiseau, FR)

Joint work of Maxim Buzdalov, Benjamin Doerr
Main reference M. Buzdalov, B. Doerr, "Runtime Analysis of the $(1 + (\lambda, \lambda))$ Genetic Algorithm on Random Satisfiable 3-CNF Formulas", Proceedings of Genetic and Evolutionary Computation Conference, 2017.
An extended version is available at arxiv: https://arxiv.org/abs/1704.04366
URL http://dx.doi.org/10.1145/3071178.3071297

The $(1+(\lambda, \lambda))$ genetic algorithm, first proposed at GECCO 2013, showed a surprisingly good performance on some optimization problems. The theoretical analysis so far was restricted to the OneMax test function, where this GA profited from the perfect fitness-distance correlation. In this work, we conduct a rigorous runtime analysis of this GA on random 3-SAT instances in the planted solution model having at least logarithmic average degree, which are known to have a weaker fitness distance correlation.

We prove that this GA with fixed not too large population size again obtains runtimes better than $\Theta(n \log n)$, which is a lower bound for most evolutionary algorithms on pseudo-Boolean problems with unique optimum. However, the self-adjusting version of the GA risks reaching population sizes at which the intermediate selection of the GA, due to the weaker fitness-distance correlation, is not able to distinguish a profitable offspring from others. We show that this problem can be overcome by equipping the self-adjusting GA with an upper limit for the population size. Apart from sparse instances, this limit can be chosen in a way that the asymptotic performance does not worsen compared to the idealistic OneMax case. Overall, this work shows that the $(1 + (\lambda, \lambda))$ GA can provably have a good performance on combinatorial search and optimization problems also in the presence of a weaker fitness-distance correlation.

3.7 Is it necessary to perform multi-objective optimization when doing multiobjectivization?

Arina Buzdalova (ITMO University – St. Petersburg, RU)

Joint work of Irina Petrova, Maxim Buzdalov, Arina Buzdalova
Main reference A. Buzdalova, I. Petrova and M. Buzdalov, "Runtime analysis of different Approaches to select conflicting auxiliary objectives in the generalized OneMax problem," 2016 IEEE Symposium Series on Computational Intelligence (SSCI), Athens, 2016, pp. 1–7
URL https://doi.org/10.1109/SSCI.2016.7850140

It has been shown that single-objective optimization may be improved by introducing auxiliary objectives. In practice, the auxiliary objectives may be conflicting. This talk presents theoretical analysis of different approaches of using auxiliary objectives on the Generalized OneMax problem with conflicting auxiliary objectives OneMax and ZeroMax.

In most of the considered methods, the optimized objectives are selected dynamically. Particularly, the $O(n \log n)$ runtime is proven for a multi-objective algorithm that optimizes the target objective together with a dynamically selected auxiliary objective. At the same time, it is shown that asymptotically the same runtime holds for a single-objective algorithm with the preservation of the best found solution, where objectives are dynamically selected using reinforcement learning.

Acknowledgements
This work was supported by RFBR according to the research project No. 16-31-00380 mol_a.

3.8 On the Variable Interaction Graph in Gray-Box Optimization

Francisco Chicano (University of Málaga, ES)

Joint work of Darrell Whitley, Renato Tinós, Gabriela Ochoa, Andrew M. Sutton, Francisco Chicano
Main reference L. Darrell Whitley, Francisco Chicano, Brian W. Goldman: Gray Box Optimization for Mk Landscapes (NK Landscapes and MAX-kSAT). Evolutionary Computation 24(3): 491-519 (2016)
URL https://doi.org/10.1162/EVCO_a_00184

Given a pseudo-Boolean function, the Variable Interaction Graph is defined using the set of variables as node set and joining two variables with an edge when there is a nonzero Walsh coefficient in the Walsh expansion of the function whose index contains both variables. In the case of k-bounded pseudo-Boolean functions (like NK Landscapes or MAX-kSAT) the co-occurrence graph can be used as an approximation of the Variable Interaction Graph. Two operators and one search algorithm based on the Variable Interaction Graph are described in the talk. We also present some open questions regarding theoretical results related to the performance of the algorithms and the optimal parameter settings based on the Variable Interaction Graph.

3.9 Fast Genetic Algorithms

Benjamin Doerr (Ecole Polytechnique – Palaiseau, FR)

For genetic algorithms using a bit-string representation of length n, the general recommendation is to take $1/n$ as mutation rate. In this work, we discuss whether this is justified for multimodal functions. Taking jump functions and the $(1+1)$ evolutionary algorithm as the simplest example, we observe that larger mutation rates give significantly better runtimes. For the $\text{Jump}_{m,n}$ function, any mutation rate between $2/n$ and m/n leads to a speed-up at least exponential in m compared to the standard choice.

The asymptotically best runtime, obtained from using the mutation rate m/n and leading to a speed-up super-exponential in m, is very sensitive to small changes of the mutation rate. Any deviation by a small $(1 \pm \varepsilon)$ factor leads to a slow-down exponential in m. Consequently, any fixed mutation rate gives strongly sub-optimal results for most jump functions.

Building on this observation, we propose to use a random mutation rate α/n, where α is chosen from a power-law distribution. We prove that the $(1+1)$ EA with this heavy-tailed mutation rate optimizes any $\text{Jump}_{m,n}$ function in a time that is only a small polynomial (in m) factor above the one stemming from the optimal rate for this m.

Our heavy-tailed mutation operator yields similar speed-ups (over the best known performance guarantees) for the vertex cover problem in bipartite graphs and the matching problem in general graphs.

Following the example of fast simulated annealing, fast evolution strategies, and fast evolutionary programming, we propose to call genetic algorithms using a heavy-tailed mutation operator *fast genetic algorithms*.

3.10 Optimal Recombination for the Asymmetric TSP: Theory and Experiment

Anton V. Eremeev (Sobolev Institute of Mathematics – Novosibirsk)

We consider two approaches to formulation and solution of optimal recombination problems arising as supplementary problems in genetic algorithms for the Asymmetric Traveling Salesman Problem. The first approach uses a representation of solutions where a genotype is the sequence of vertices visited in the traveling salesman tour (position-based encoding). The second approach uses a representation where each gene defines an arc of the tour (adjacency-based representation). Both optimal recombination problems under consideration are NP-hard but relatively fast worst-case exponential-time algorithms are presented for solving them [2]. Besides that, in the case of position-based encoding, the optimal recombination problem is shown to be solvable in linear time for "almost all" pairs of parent solutions.

As a proof of concept we develop a genetic algorithm with a crossover operator which solves an optimal recombination problem. The algorithm also incorporates problem-specific mutation operator, local search and initialization method. A computational experiment on TSPLIB instances shows that the proposed genetic algorithm yields competitive results to other state-of-the-art genetic algorithms [1].

The research is supported by Russian Science Foundation grant 15-11-10009.

References

1 Anton V. Eremeev and Yulia V. Kovalenko.*Genetic algorithm with optimal recombination for the asymmetric travelling salesman problem*. To appear in Pros. of LSSC-17, Sozopol, Bulgaria, 2017

2 Anton V. Eremeev and Yulia V. Kovalenko.*Experimental evaluation of two approaches to optimal recombination for permutation problems*. In Proc. of EvoCOP-16, Porto, Portugal, Springer, Switzerland, LNCS vol. 9595. pp. 138–153 (2016)

3.11 Mathematical Models of Artificial Genetic Representations with Neutrality

Carlos M. Fonseca (University of Coimbra, PT)

Joint work of Carlos M. Fonseca, Vida Vukašinović, Nino Bašić

In this work, a mathematical framework for the study and characterisation of a family of uniformly redundant binary representations based on error-control codes proposed previously [1] is developed. Such representations can exhibit various degrees of redundancy, neutrality, and other properties believed to influence the performance of evolutionary algorithms, such as connectivity, locality, and synonymity [2], and have allowed this influence to be studied experimentally to some extent [3]. The definition of suitable equivalence classes leads to a partitioning of the representation space with respect to neutral network structure and connectivity, which should allow the effect of locality on search performance to be studied while other properties are kept fixed. The practical implications of the proposed framework are also discussed.

Acknowledgements. This talk is based upon work from COST Action CA15140 on Improving Applicability of Nature-Inspired Optimisation by Joining Theory and Practice (ImAppNIO), supported by COST (European Cooperation in Science and Technology). Partial support by national funds through the Portuguese Foundation for Science and Technology (FCT) and by the European Regional Development Fund (FEDER) through COMPETE 2020 – Operational Programme for Competitiveness and Internationalisation (POCI) is also acknowledged.

References

1 C. M. Fonseca and M. B. Correia. "Developing redundant binary representations for genetic search," in *The 2005 IEEE Congress on Evolutionary Computation*, vol. 2, (Edinburgh, U.K.), pp. 1675–1682, Sept. 2005.
2 F. Rothlauf. *Representations for Genetic and Evolutionary Algorithms*. Springer, 2nd Edition, 2006.
3 M. B. Correia. "A study of redundancy and neutrality in evolutionary optimization," *Evolutionary Computation*, vol. 21, pp. 413–443, 2013.

3.12 Monotone Functions on Bitstrings – Some Structural Notes

Christian Gießen (Technical University of Denmark – Lyngby, DK)

Joint work of Tobias Friedrich, Timo Kötzing, Martin Schirneck, Christian Gießen

Analyzing the (1+1) EA with mutation probability c/n on monotone functions for a constant c is a challenging problem. Only for $c \leq 1$ and $c > 2.2$ upper bounds on the expected runtime have been established, but the whole truth is unkown. We present unfinished work that led to two curious and surprising structural observations. First, the linear function BinVal is a special monotone function: globally, it is an optimal adversarial function in the sense that it maximizes the possibility to increase the distance to the optimum while increasing the fitness at the same time. This result can be shown using techniques that stem from extremal set theory and general order theory and is related to the Shadow Minimization Problem. However, this notion of hardness is not sufficient to describe hardness for the (1+1) EA. The second surprising observation is the distribution of monotone functions that are structurally equivalent for small n. OneMax-like functions occur most often, while BinVal-like functions occur the least. It is however unclear if this behaviour holds for all n. It remains an open question how to structurally define hardness for monotone functions.

References

1 The on-line encyclopedia of integer sequences. published electronically at https://oeis.org. https://oeis.org/A000372.

2 The on-line encyclopedia of integer sequences. published electronically at https://oeis.org. https://oeis.org/A046873.
3 Graham R. Brightwell and Prasad Tetali. The number of linear extensions of the boolean lattice. *Order*, 20(4):333–345, 2003.
4 Sylvain Colin, Benjamin Doerr, and Gaspard Férey. Monotonic functions in EC: anything but monotone! In *Genetic and Evolutionary Computation Conference, GECCO '14, Vancouver, BC, Canada, July 12-16, 2014*, pages 753–760, 2014.
5 Benjamin Doerr, Thomas Jansen, Dirk Sudholt, Carola Winzen, and Christine Zarges. Mutation rate matters even when optimizing monotonic functions. *Evolutionary Computation*, 21(1):1–27, 2013.
6 K. Engel. *Sperner Theory.* Cambridge Solid State Science Series. Cambridge University Press, 1997.
7 S.R. Finch. *Mathematical Constants.* Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2003.
8 Graham. The number of linear extensions of ranked posets. Technical report, London School of Economics, 2003.
9 Thomas Jansen. On the brittleness of evolutionary algorithms. In *FOGA*, pages 54–69, 2007.
10 Jichang Sha and Daniel J. Kleitman. The number of linear extensions of subset ordering. *Discrete Mathematics*, 63(2-3):271–278, 1987.

3.13 Global Convergence of the (1+1)-ES

Tobias Glasmachers (Ruhr-Universität Bochum, DE)

We prove several global convergence theorems for the (1+1)-ES algorithm. We refrain from inserting nursing mechanisms into the algorithms, and at no point we resort to asymptotic or otherwise approximate analysis. Therefore the theorems and their proofs reflect the actual behavior of the algorithm.

The analysis is based on two ingredients. We start with a generic sufficient decrease condition for elitist rank-based evolution strategies, formulated for an essentially monotonically transformed variant of the objective function. Then we show that the algorithm state is found infinitely often in a regime where step size and success rate are simultaneously bounded away from zero, with full probability.

The main result is proven by combining these statements. More powerful variants are derived based on additional regularity conditions. The statements ensure under minimal technical preconditions that the sequence of iterates has a limit point in a critical point of some sort.

Based on our theorems we analyze the behavior of the (1+1)-ES on a number of problems ranging from the smooth (non-convex) cases over various types of ridge functions to several discontinuous problems.

3.14 How to Guaranty Positive Definiteness in Active CMA-ES

Nikolaus Hansen (INRIA Saclay – Île-de-France, FR)

Main reference Nikolaus Hansen, The CMA Evolution Strategy: A Tutorial. CoRR abs/1604.00772 (2016)
URL http://arxiv.org/abs/1604.00772

In active CMA-ES, a positive definite covariance matrix undergoes an additive update using also negative weights. Depending on the chosen weights and learning rate, the covariance matrix may become negative definite with a small probability. We can prove a condition on the weights which guaranties positive definiteness with probability one, if also the length of the random vectors is upper bounded. The condition limits the decrement of the smallest eigenvalue to a factor of about one over dimension in each iteration. We investigate the population size when the constraint becomes effective as a function of dimension.

3.15 An overview of Bayesian Optimization

Matthew W. Hoffman (Google DeepMind – London, GB)

Joint work of Bobak Shariari, Matthew W. Hoffman
Main reference Matthew W. Hoffman, Zoubin Ghahramani. "Output-Space Predictive Entropy Search for Flexible Global Optimization." NIPS Workshop on Bayesian Optimization, 2015.

In this talk I give a high-level overview of Bayesian optimization methods for global optimization. While typically used for continuous, black-box optimization I also briefly touch on relations to the broader optimization community.

Further, while the design of methods for Bayesian optimization involves a great number of choices that are often implicit in the overall algorithm design, in this work I argue for a modular approach to Bayesian optimization. In particular this includes selection of the acquisition function, kernel, and hyper-priors as well as less-discussed components such as the recommendation and initialization approaches. In this work I also argue for an information-theoretic approach to the design of acquisition strategies. Finally, in this work I present a Python implementation, pybo, that allows us to easily vary these choices. Ultimately this approach provides us a straightforward mechanism to examine the effect of each choice both individually and in combination.

3.16 Estimation of Distribution Algorithms

Martin S. Krejca (Hasso-Plattner-Institut – Potsdam, DE)

Joint work of Tobias Friedrich, Timo Kötzing, Andrew M. Sutton, Martin S. Krejca

Evolutionary algorithms (EAs) are optimization techniques inspired by nature. They are a popular choice if the problem at hand is, for example, noisy or highly complex and cannot be well formalized but the quality of a single solution can be easily measured. Typically, EAs maintain a set of samples from the solution space, which is iteratively updated, keeping better solutions and discarding bad ones. An alternative and more direct approach that is also commonly used is to not store samples but a probability distribution over the search space that generates these samples. Such algorithms are called estimation of distribution algorithms (EDAs).

In practice, EDAs are widely applied and perform very well. However, theoretical results on EDAs explaining this success are very scarce so far. We introduce an EDA framework we proposed, which subsumes many EDAs used for discrete domains, and we present our theoretical results for this framework. This includes robustness of EDAs to noise, restrictions on the way an EDA can update its distribution, and unbiasedness.

3.17 Convergence in Genetic Programming

William B. Langdon (University College London, GB)

Main reference William B. Langdon, "Long-Term Evolution of Genetic Programming Populations," GECCO 2017: The Genetic and Evolutionary Computation Conference", 2017.
URL https://arxiv.org/abs/1703.08481

In the future we hope to use evolution to solve challangeing problems. This may require it to be run for many generations. Therefore we investigate what happens in long runs of a current evolutionary algorithms. Specifically we look at an easy problem 6-mux when solved by genetic programming when the GP system is run on long past the point where GP solves the six multiplexor problem.

We evolve binary mux-6 trees for up to 100000 generations evolving some programs with more than a hundred million nodes. Initially tree growth is O(generations squared), [GP+EM (1)1 pp95-119] but existing theory could be made more formal. Long after the first time when everyone in the finit population solves the problem, and so has identical fitness, the tree size *appears* to execute a random walk, albeit with a lower bound. However, even in this region, the distribution of tree sizes within the population is not a gammar distribution as predicted for large populations and no selection (http://www.cs.bham.ac.uk/~wbl/biblio/gp-html/poli_2007_eurogp.html). Our unbounded Long-Term Evolution Experiment LTEE GP appears not to evolve building blocks but does suggests a limit to bloat. We do see periods of tens even hundreds of generations where the population is 100 percent functionally converged.

3.18 Landscape of the Triangle Program

William B. Langdon (University College London, GB)

Joint work of William B. Langdon, Nadarajen Veerapen, Gabriela Ochoa
Main reference William B. Langdon, Nadarajen Veerapen, Gabriela Ochoa, "Visualising the Search Landscape of the Triangle Program," Genetic Programming – 20th European Conference, 96–113, 2017.
URL http://dx.doi.org/doi:10.1007/978-3-319-55696-3_7

the triangle program is a small software engineering benchmark recently analysed in terms of global search, genetic algorithms schema, iterated local search and local optima networks. Results presented at EuroGP-2017 (http://dx.doi.org/doi:10.1007/978-3-319-55696-3_7).

I build on these results to support the thesis that real software is not as fragile as is often assumed and then consider in detail the simplest of the test cases – the scalene triangle, looking at the variable interaction graph for the test case when only comparison operators are to be mutated. I propose the variable interaction graph in real software may lead to theoretical insights to the improvement of sizable programs using evolutionary improvement methods such as genetic programming.

The triangle benchmark is available via http://www.cs.ucl.ac.uk/staff/w.langdon/gggp/#triangle_dataset.

3.19 Runtime Analysis Evolutionary Algorithms with Self-adaptive Mutation Rates

Per Kristian Lehre (University of Birmingham, GB)

Joint work of Duc-Cuong Dang, Per Kristian Lehre
Main reference Duc-Cuong Dang, Per Kristian Lehre, "Self-adaptation of Mutation Rates in Non-elitist Populations," Parallel Problem Solving from Nature – PPSN XIV, 803–813, 2016.
URL http://www.cs.nott.ac.uk/~pszpl/selfadapt/

The runtime of evolutionary algorithms (EAs) depends critically on their parameter settings, which are often problem-specific. Automated schemes for parameter tuning have been developed to alleviate the high costs of manual parameter tuning.

Experimental results indicate that self-adaptation, where parameter settings are encoded in the genomes of individuals, can be effective in continuous optimisation. However, results in discrete optimisation have been less conclusive. Furthermore, a rigorous runtime analysis that explains how self-adaptation can lead to asymptotic speedups has been missing.

This talk presents the first runtime analysis of self-adaptation for discrete, population-based EAs. We apply the level-based theorem to show how a self-adaptive EA is capable of fine-tuning its mutation rate, leading to exponential speedups over EAs using fixed mutation rates.

For a simulation and a link to the paper, please see http://www.cs.nott.ac.uk/~pszpl/selfadapt/.

3.20 Noise models for comparison-based evolutionary algorithms

Johannes Lengler (ETH Zürich, CH)

Noise models for evolutionary algorithms typically assume that the fitness function evaluation is distorted. However, for comparison-based evolutionary algorithms this makes only limited sense, since often such algorithm do not actually evaluate a fitness function. I have discussed three examples where fitness functions are not evaluated, and their implications for noise models: swap-based sorting, Schöning's algorithms, and the evolution of game engines.

References

1 Tomáš Gavenčiak, Barbara Geissmann, and Johannes Lengler. *Sorting by Swaps with Noisy Comparisons*. Proceedings of the 2017 conference on Genetic and Evolutionary Computation (GECCO '17)

3.21 Features, Diversity, Random Walks and Digital Art

Frank Neumann (University of Adelaide, AU)

Joint work of Frank Neumann, Wanru Gao, Samadhi Nallaperuma, Aneta Neumann, Bradley Alexander, James Kortman, Zygmunt Szpak, Wojciech Chojnacki

We consider diversity with respect to features of a given problem and introduce an evolutionary algorithm that maximizes the feature diversity under a constraint given by a lower bound on the desired objective function value (in the case of maximization) [1]. We show how this increases feature diversity for the problem of evolving hard TSP instances (in terms of approximation ratio) for the classical 2-opt local search algorithm. Furthermore, we discuss how this feature-based diversity approach can be used to create variation of images (based on a given image) with respect to various features [2]. Afterwards, we introduce evolutionary approaches for image transition and composition based on random walks [3]. The evolutionary image transition approach allows to create artistic sequences of images in form of a video by transforming a given starting image into a target image. Our evolutionary image composition approach combines two artistic images taken into account interesting parts of the given images [4]. The key element is a fitness function based on covariance matrix image descriptors taking into account a set of features.

References

1 W. Gao, T. Friedrich, F. Neumann: Fixed-parameter single objective search heuristics for minimum vertex cover. In: Parallel Problem Solving from Nature XIII, PPSN 2016, Springer, 2016.

2 B. Alexander, J. Kortman, A. Neumann: Evolution of artistic image variants through feature based diversity optimisation. In: Genetic and Evolutionary Computation Conference, GECCO 2017, ACM Press, 2017.

3 A. Neumann, B. Alexander, F. Neumann: Evolutionary image transition using random walks. In: International Conference on Computational Intelligence in Music, Sound, Art and Design, EVOMUSART 2017, Springer, 2017.

4 A. Neumann, Z. L. Szpak, W. Chojnacki, F. Neumann: Evolutionary image composition using feature covariance matrices. In: Genetic and Evolutionary Computation Conference, GECCO 2017, ACM Press, 2017.

3.22 Standard Steady State Genetic Algorithms can Hillclimb Faster than Mutation-only Evolutionary Algorithms

Pietro S. Oliveto (University of Sheffield, GB)

Joint work of Dogan Corus, Pietro S. Oliveto

Explaining to what extent the real power of genetic algorithms lies in the ability of crossover to recombine individuals into higher quality solutions is an important problem in evolutionary computation. Recently it has been shown how the interplay of mutation and crossover may create the necessary diversity to efficiently escape local optima. In this talk we discuss how such an interplay can also make genetic algorithms hillclimb faster than their mutation-only counterparts. We devise a Markov Chain framework that allows to rigorously prove an upper bound on the runtime of standard steady state genetic algorithms to hillclimb the ONEMAX function. The bound establishes that the steady-state genetic algorithms are 25% faster than all unbiased mutation-only evolutionary algorithms with static mutation rate up to lower order terms for moderate population sizes. The analysis also suggests that larger populations may be faster than populations of size 2. We present a lower bound for a greedy (2+1) GA that matches the upper bound for populations larger than 2, rigorously proving that 2 individuals cannot outperform larger population sizes under greedy selection and greedy crossover up to lower order terms. In complementary experiments the best population size is greater than 2 and the greedy genetic algorithms are faster than standard ones, further suggesting that the derived lower bound also holds for the standard steady state (2+1) GA.

3.23 Linear multiobjective drift analysis

Jonathan E. Rowe (University of Birmingham, GB)

Drift analysis is now a standard tool for analysing the run-time of stochastic optimisation algorithms. The expected progress in the "distance" to the target state is used to derive a first hitting time of that state. We consider the situation where the process is best described by more than distance function, and derive a generalisation of multiplicative drift to this situation. Example applications include: an evolutionary algorithm solving a multi-objective optimisation problem; a parallel island model with probabilistic migration.

3.24 Theoretical Aspects of the Averaged Hausdorff Indicator in Biobjective Optimization

Günter Rudolph (TU Dortmund, DE)

Joint work of Oliver Schütze, Heike Trautmann, Günter Rudolph
Main reference Günter Rudolph, Oliver Schütze, Christian Grimme, Christian Domínguez-Medina, Heike Trautmann, "Optimal averaged Hausdorff archives for bi-objective problems: theoretical and numerical results." Comp. Opt. and Appl., 64(2): 589-618, 2016.

The averaged Hausdorff indicator is an alternative to the dominated hypervolume indicator for assessing the quality of Pareto front approximations if equispaced solutions on the Pareto front are desired. This may happen frequently in dynamic control applications. Recently there have been several proposals how to integrate this indicator in evolutionary search for up to 4 objectives. Empirical evaluations have shown that this approach is promising. Therefore a theoretical analysis of the indicator is desireable. We show results in case of two objectives. Point sets that have minimal averaged Hausdorff distance to the Pareto front are called optimal archives. If the Pareto front is concave, then the optimal archives are on the Pareto front. If the Pareto front is linear, the optimal archive consists of equispaced solutions on the Pareto front. If the Pareto front is circularly concave, then the optimal archive consists of equispaced points. If the Pareto front is circularly concave, then the optimal archive consists of equispaced points but no element is on the Pareto front. But it can be shown that the averaged Hausdorff distance of the archive to the Pareto front decreases to zero with order $1/m^2$ for increasing archive size m. For practical purposes with archive sizes $m \geq 100$ the accuracy, i.e. the closeness to the Pareto front, is sufficient.

3.25 Max-Min Thompson Sampling for the K-Armed Dueling Bandit Problem

Jonathan L. Shapiro (University of Manchester, GB) and Joseph Mellor

Joint work of Joseph Mellor, Jonathan Shapiro

The K-armed dueling bandit problem is a variation on the classic K-armed bandit problem. The distinguishing feature of the problem is that only relative preference between pairs of arms is given as feedback. This paper proposes a new algorithm, Max-Min Thompson Sampling, to solve the problem. The algorithm uses a method derived from game theory to choose appropriate pairs of arms, and Thompson Sampling to learn the preferences from observations. We derive an $O(K \log T)$ problem-dependent and finite-time regret bound for the strategy, where T is the time. Our bound is as low as others in the literature. We provide empirical results of the method on a variety of simulations including the Komiyama et.al. (2015) benchmarks, and a real-world information retrieval task. These results show very strong performance on the simulations investigated. The use of game theory as a principle suggests other applications and extensions of the method.

References

1 Komiyama, Junpei and Honda, Junya and Kashima, Hisashi and Nakagawa, Hiroshi *Regret Lower Bound and Optimal Algorithm in Dueling Bandit Problem*. COLT 2015,

http://jmlr.org/proceedings/papers/v40/Komiyama15.html, JMLR Proceedings, Vol. 40, pp 1141-1154, 2015.

3.26 Fundamentals of ESs' Statistical Learning

Ofer M. Shir (Tel-Hai College – Upper Galilee, IL)

Joint work of Ofer M. Shir, A. Yehudayoff
Main reference Shir, O. M. and Yehudayoff, A. (2017). On the statistical learning ability of evolution strategies. In Proceedings of the 14th ACM/SIGEVO Conference on Foundations of Genetic Algorithms, FOGA'17, pages 127–138, New York, NY, USA. ACM.
URL http://dx.doi.org/10.1145/3040718.3040722

We consider Evolution Strategies operating only with isotropic Gaussian mutations on positive quadratic objective functions and investigate the Covariance matrix when constructed out of selected individuals by truncation. We prove that the statistically constructed Covariance matrix over such selected decision vectors becomes proportional to the inverse of the landscape Hessian as the population size increases. This generalizes a previous result that proved an equivalent phenomenon when sampling is carried out in the vicinity of the optimum [FOGA'17]. It further confirms the classical hypothesis that statistical learning of the landscape is an inherent characteristic of standard ESs, and that this distinguishing capability stems only from the usage of isotropic Gaussian mutations and rank-based selection.

References

1 Shir, O. M. and Yehudayoff, A. On the statistical learning ability of evolution strategies. In Proceedings of the 14th ACM/SIGEVO Conference on Foundations of Genetic Algorithms, FOGA'17, pages 127–138, New York, NY, USA. ACM, 2017

3.27 Low discrepancy for one-shot optimization

Olivier Teytaud (Google Switzerland – Zürich, CH)

Joint work of Google Brain, Zurich

The selection of hyper-parameters is critical in Deep Learning. Because of the long training time of complex models and the availability of compute resources in the cloud, "one-shot" optimization schemes – where the sets of hyperparameters are selected in advance (either on a grid or in a random manner) and the training is executed in parallel – are commonly used. (Bergstra & Bengio, 2012) show that grid search is sub-optimal, especially when only a few parameters matter, and suggest to use random search instead. Yet, random search can be "unlucky" and produce sets of values that leave some part of the domain unexplored. Quasi-random methods, such as Low Discrepancy Sequences (LDS) avoid these issues. We show that such methods have theoretical properties that make them appealing for performing hyperparameter search, and demonstrate that, when applied to the selection of hyperparameters of complex Deep Learning models (such as state-of-the-art LSTM language models), they yield suitable hyperparameters values with much fewer runs than random search. We propose a particularly simple LDS method which can be used as a drop-in replacement for grid/random search in any Deep Learning pipeline.

3.28 Recent Advances in Runtime Analysis of Estimation-of-Distribution Algorithms

Carsten Witt (Technical University of Denmark – Lyngby, DK)

We consider three simple estimation-of-distribution algorithms (EDAs) on the OneMax benchmark function. The runtime of these algorithms depends on the number of bits n and the precision λ of the model it builds. Exponential runtimes as well as polynomial runtimes of the kind $\Theta(\lambda n)$ and $\Theta(\lambda\sqrt{n})$ are obtained in different regimes for λ. Two phase transitions in the runtime behavior are identified, and open problems are discussed.

Literature:

References

1 Tianshi Chen, Ke Tang, Guoliang Chen, and Xin Yao. On the analysis of average time complexity of estimation of distribution algorithms. In *Proc. of CEC '07*, pages 453–460, 2007.
2 Tianshi Chen, Per Kristian Lehre, Ke Tang, and Xin Yao. When is an estimation of distribution algorithm better than an evolutionary algorithm? In *Proc. of CEC '09*, pages 1470–1477, 2009a.
3 Tianshi Chen, Ke Tang, Guoliang Chen, and Xin Yao. Rigorous time complexity analysis of univariate marginal distribution algorithm with margins. In *Proc. of CEC '09*, pages 2157–2164, 2009b.
4 Tianshi Chen, Ke Tang, Guoliang Chen, and Xin Yao. Analysis of computational time of simple estimation of distribution algorithms. *IEEE Transactions on Evolutionary Computation*, 14 (1): 1–22, 2010.
5 Duc-Cuong Dang and Per Kristian Lehre. Simplified runtime analysis of estimation of distribution algorithms. In *Proc. of GECCO '15*, pages 513–518, 2015.
6 Stefan Droste. A rigorous analysis of the compact genetic algorithm for linear functions. *Natural Computing*, 5(3): 257–283, 2006.
7 Tobias Friedrich, Timo Kötzing, Martin S. Krejca, and Andrew M. Sutton. The benefit of recombination in noisy evolutionary search. In *Proc. of ISAAC '15*, pages 140–150, 2015.
8 Tobias Friedrich, Timo Kötzing, and Martin S. Krejca. EDAs cannot be balanced and stable. In *Proc. of GECCO '16*, pages 1139–1146, 2016.
9 Martin S. Krejca and Carsten Witt. Lower bounds on the run time of the univariate marginal distribution algorithm on OneMax. In *Proc. of FOGA 2017*, pages 65–79, 2017.
10 Per Kristian Lehre and Phan Trung Hai Nguyen. Improved runtime bounds for the univariate marginal distribution algorithm via anti-concentration. In *Proc. of GECCO '17*, 2017. To appear.
11 Frank Neumann, Dirk Sudholt, and Carsten Witt. A few ants are enough: ACO with iteration-best update. In *Proc. of GECCO '10*, pages 63–70, 2010.
12 Dirk Sudholt and Carsten Witt. Update strength in EDAs and ACO: How to avoid genetic drift. In *Proc. of GECCO '16*, pages 61–68, 2016.

13 Carsten Witt. Upper Bounds on the Runtime of the Univariate Marginal Distribution Algorithm on OneMax. In *Proc. of GECCO '17*, 2017. To appear. Preprint: http://arxiv.org/abs/1704.00026

3.29 (1+λ) Evolutionary Algorithm with Self-Adjusting Mutation Rate

Jing Yang (Ecole Polytechnique – Palaiseau, FR)

Joint work of Benjamin Doerr, Christian Gießen, Carsten Witt, Jing Yang

We propose a self-adapting strategy for the population-based evolutionary algorithms. It creates half the offspring with a mutation rate that is twice the current mutation rate and the other half with half the current rate. The mutation rate is then updated to the rate which generates the best offspring. We prove that the self-adapting (1+λ) EA solves OneMax in an expected generations of $O(n/\log(\lambda) + n\log(n)/\lambda)$. According to previous work, this is best-possible among all lambda-parallel mutation-based unbiased black-box algorithms.

4 Working groups

4.1 Breakout Session: Information Geometric Optimization

Youhei Akimoto (Shinshu University – Nagano, JP)

Joint work of Youhei Akimoto, Dirk Arnold, Anne Auger, Hans-Georg Beyer, Dimo Brockhoff, Tobias Glasmachers, Nikolaus Hansen, Christian Igel

Information geometric optimization (IGO) [1] is a generic framework for probability model based search algorithms for arbitrary search space. Given a family of probability distributions over the search space, a probability model based search algorithm like an estimation of distribution algorithm (EDA) is derived. The population based incremental learning (PBIL) for binary optimization, compact genetic algorithm (cGA), and some components of the covariance matrix adaptation evolution strategy (CMA-ES) can be derivedfrom the IGO framework, while they have been proposed independently from it. The development of the IGO framework has so far contributed to analyze a simplified CMA-ES [2] and to derive a novel variant of the CMA-ES [3]. However, as the outcome of these results, we find that the analyzed model of the CMA-ES behaves differently from what we observe in practice. Moreover, the components of the CMA-ES that are not included in the IGO framework such as the step-size adaptation and the rank-one covariance matrix update are often critical to the performance, hence we need to incorporate them when deriving a novel algorithm. The objective of the breakout session was to share our understanding of the algorithms and the IGO framework and to develop the framework so that it provides more reasonable mathematical models for analysis and more practical algorithms. In this breakout session, we started with discussing how the rank-one covariance matrix update can be interpreted. Then, we discussed why the ordinary differential equation that the IGO framework provides as a mathematical model to analyze behaves differently from the behavior of the real algorithm, and how we can correct it. The use of stochastic differential equations was proposed. In relation to this, we also discussed the dependency of the algorithm performance on the choice of the parameters of the probability distribution. Finally, we shared recent developments on the IGO and related algorithms such as the population size adaptation [4].

References

1 Yann Ollivier, Ludovic Arnold, Anne Auger, Nikolaus Hansen. Information-Geometric Optimization Algorithms: A Unifying Picture via Invariance Principles, Journal of Machine Learning Research 18(18), pages 1-65, 2017.
2 Hans-Georg Beyer. Convergence Analysis of Evolutionary Algorithms That are Based on the Paradigm of Information Geometry, Evolutionary Computation 22(4), pages 679-709, 2014.
3 Youhei Akimoto, Anne Auger, Nikolaus Hansen. Comparison-Based Natural Gradient Optimization in High Dimension, In Proceedings of the Genetic and Evolutionary Computation Conference, pages 373-380, 2014.
4 Kouhei Nishida, Youhei Akimoto. Population Size Adaptation for the CMA-ES Based on the Estimation Accuracy of the Natural Gradient, In Proceedings of the Genetic and Evolutionary Computation Conference, pages 237-244, 2016.

4.2 Breakout Session: Preference-based Selection in Evolutionary Multiobjective Optimisation

Carlos M. Fonseca (University of Coimbra, PT)

Participants: Dimo Brockoff, Arina Buzdalova, Francisco Chicano, Carlos M. Fonseca, Jonathan L. Shapiro

Multiobjective optimisation consists in simultaneously optimising two or more incommensurable objective functions over the same domain. In practice, however, the optima of all such objective functions seldom coincide, and there is no ideal solution. Instead, there are usually multiple incomparable solutions that are optimal (or *efficient*) in the sense that no other solution is at least as good in all objectives and strictly better in at least one of them. Therefore, a compromise solution is often sought among all efficient solutions. The image of the set of efficient solutions in objective space is known as the Pareto front.

Selecting a single compromise solution involves a decision making process, and requires additional information, known as preference information, which may not be explicitly available when optimisation is performed. One approach to this situation consists in searching for a diverse set of efficient solutions, in the hope that it will contain a suitable compromise solution. Traditionally, preference information has been considered to pertain to individual candidate solutions, but this notion has meanwhile been extended to sets of candidate solutions, including the preference for diverse sets of solutions.

The aim of this breakout session was to discuss how solution-oriented preferences and set-oriented preferences relate to each other, and to what extent they can be meaningfully combined. The starting point for the discussion was the observation that diversity appears to be in contradiction with solution-oriented preferences. In particular, if the preferences of the Decision Maker (DM) are fully known in advance, why should optimisation provide a diverse set of non-preferred solutions?

In Evolutionary Multiobjective Optimisation (EMO), three main approaches to diversity promotion can be identified. Initially, selection was guided by the solution-oriented preferences, and diversity was promoted by other means, such as niching techniques [1]. In indicator-based approaches [2] and decomposition-based approaches [3], on the other hand, diversity is promoted directly via selection, but in different ways. Whereas quality indicators directly specify what a good set is, in decomposition-based approaches multiple solution-oriented preferences are considered, leading to multiple preferred solutions. In can be argued that both of these approaches attempt to approximate the Pareto front in some way, where approximation quality is tied to set-oriented preferences that must be expressed by the DM.

An alternative view embraces the fact that DM preferences are seldom well understood in advance, and adopts a probabilistic view of the associated uncertainty. As a result, the quality of a single solution is itself a random variable, and selection can be reinterpreted as a portfolio optimisation problem [4]. This entails modelling the uncertainty about how good each solution is and the dependence between the quality of pairs of solutions, rather than expressing preferences about sets. Solving the portfolio optimisation problem consists of maximising expected return (i.e., some measure of DM satisfaction) and balancing it against deviations from this expectation (risk), leading to diverse solutions in the portfolio. In contrast to other approaches, diversity emerges as a risk-balancing strategy due to preference uncertainty, rather than being preferred as such.

From the discussion, the formulation of alternative preference uncertainty models, ideally supporting DM interaction, the study of alternative portfolio optimisation formulations leading to simpler computational problems (e.g., linear versus quadratic programming) and the characterisation of the resulting quality indicators were identified as current research challenges.

References

1. C. M. Fonseca and P. J. Fleming. "Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization," in *Genetic Algorithms: Proceedings of the Fifth International Conference*, pp. 416–423, San Mateo, CA: Morgan Kaufmann, 1993.
2. E. Zitzler and S. Künzli. "Indicator-based selection in multiobjective search," in *Parallel Problem Solving from Nature – PPSN VIII*, vol. 3242 of LNCS, pp. 832–842, Springer, 2004.
3. Q. Zhang and H. Li. "MOEA/D: A multiobjective evolutionary algorithm based on decomposition," *IEEE Transactions on Evolutionary Computation*, vol. 11, pp. 712–731, 2007.
4. I. Yevseyeva, A. P. Guerreiro, M. T. M. Emmerich, and C. M. Fonseca. "A portfolio optimization approach to selection in multiobjective evolutionary algorithms." in *Parallel Problem Solving from Nature – PPSN XIII*, vol. 8672 of LNCS, pp. 672–681, Springer, 2014.

4.3 Breakout Session: Drift Theorems for Continuous Optimization

Tobias Glasmachers (Ruhr-Universität Bochum, DE)

Participants: Nikolaus Hansen, Anne Auger, Per Kristian Lehre, Günter Rudolph, Johannes Lengler, Carlos Fonseca, Youhei Akimoto, Carsten Witt, Hans-Georg Beyer, Tobias Glasmachers

Concept of the session:

- inform everyone on current state-of-the-art drift theorems
- collect requirements to make drift analysis applicable to optimization and online parameter control in continuous optimization
- see what can already be done, check discrepancies, check what can be improved and adapted in the theorems to make them more directly applicable to continuous problems

We started with a short presentation given by Per Kristian. It turns out that even very classic results are directly applicable to continuous problems. The seeming discrepancy was mostly caused by specific adaptations of general drift theorems (coming from supermartingale theory) to the specific needs of discrete optimization. The only relevant limitation seems be that spaces need to be bounded in order to obtain lower bounds, but expected runtime and upper bounds can be derived in surprising generality. Johannes took over, explaining a simple and insightful proof, demonstrating why prerequisites are more strict for lower bounds.

Results are generally of the following form. Given a stochastic process $(Y_k)_{k\in\mathbb{N}}$ and a drift condition $E[Y_{k+1}|Y_k] \leq d$, the expected runtime for hitting a lower bound a is simply $(Y_0 - a)/d$. Statements for quantiles look similar and hold with overwhelming probability (one minus exponentially small exception). Besides this additive drift there is also multiplicative drift and variable drift.

These results seem to be generally well suited for the analysis of evolution strategies, for all state variables (optimization progress, step size and covariance matrix adaptation, and evolution paths). However, finding good potential functions may be challenging. Combined drift (presented earlier by Jonathan Rowe) seems to be highly relevant.

As proposed by Anne, the next step is to try this out in simple instructive cases, like the 1/5 success rule.

We closed with the question of good resources presenting these drift theorems in the form of a review paper or a concise collection. Per Kristian and Carsten will update their arXiv papers to reflect the state-of-the-art.

It turns out that Jens Jägerküpper had already used such techniques in his convergence analysis of the (1+1)-ES, although it was not recognized at the time that the technique is of very generic value, and it was not called drift.

The session went surprisingly smoothly. In contrast to the expectations, we did not encounter any serious obstacles. It seems that all prerequisites are in place for applying drift to the analysis of evolution strategies in the near future.

4.4 Breakout Session: Discrete Estimation of Distribution Algorithms

Martin S. Krejca (Hasso-Plattner-Institut – Potsdam, DE)

Participants: Youhei Akimoto, Tobias Friedrich, Christian Gießen, Nikolaus Hansen, Martin S. Krejca, Per Kristian Lehre, Johannes Lengler, Frank Neumann, Dirk Sudholt, Andrew M. Sutton, Carsten Witt

This breakout session was held to discuss the next steps to take in the analysis of estimation of distribution algorithms (EDAs). Many of the active researchers in that area were present.

We started by discussing open problems and searching for topics that everyone was interested in. The discussed topics included optimization under constraints, more general run time analyses, and dependencies. Our group was most interested in dependencies. One of the selling points of EDAs generally is that they can adapt very well to the underlying structure of a problem, which often includes dependencies. Since most of the EDAs considered so far in theory use a univariate model (assuming independence of the problem variables) [1, 2, 3, 4], it seemed important and necessary to also analyze multivariate EDAs.

After discussing different multivariate EDAs, our group agreed to start with MIMIC [5]: a bivariate EDA whose probabilistic model generates a permutation of the problem variables. A natural first problem to analyze for such an algorithm is LeadingOnes, which is the problem of finding a hidden permutation. However, we were not convinced that MIMIC will outperform univariate EDAs on that function. This lead to a discussion of MIMIC's performance on other functions like OneMax. Overall, we decided to first run some experiments to see how MIMIC performs on standard benchmarks used in theory. After that, we want to prove the observed behavior.

References

1 T. Friedrich, T. Kötzing, and M. S. Krejca. EDAs Cannot Be Balanced and Stable. In *Proc. of GECCO '16*, pp. 1139–1146, 2016.
2 P. K. Lehre and P. T. H. Nguyen. Tight Bounds on Runtime of the Univariate Marginal Distribution Algorithm via Anti-Concentration. In *Proc. of GECCO '17*, 2017.
3 D. Sudholt and C. Witt. Update Strength in EDAs and ACO: How to Avoid Genetic Drift. In *Proc. of GECCO '16*, pp. 61–68, 2016.
4 C. Witt. Upper Bounds on the Runtime of the Univariate Marginal Distribution Algorithm on OneMax. In *Proc. of GECCO '17*, 2017.
5 J. S. De Bonet, C. L. Isbell, and P. Viola. MIMIC: Finding Optima by Estimating Probability Densities. In *Proc. of NIPS '97*, pp. 424–431, 1997.

4.5 Breakout Session: Theory of Evolutionary Algorithms for Problems With (Dynamic) Constraints

Frank Neumann (University of Adelaide, AU)

Participants: Tobias Friedrich Carola Doerr Andrew M. Sutton William B. Langdon Dirk Sudholt Jing Yang Francisco Chicano Arina Buzdalova Thomas Jansen Martin Krejca Christine Zarges Christian Giessen Anton Eremeev Pietro Oliveto Dimo Brockhoff Frank Neumann

Most of the results in the area of runtime analysis are for unconstrained problems. However, in practice constraints play a crucial role. The aim of the breakout session has been to explore new research directions to provide meaningful theoretical insights into the working behavior of evolutionary computing techniques for problems with dynamically changing constraints. Starting the discussion, it has been observed that understanding evolutionary computing techniques for problems with constraints is a rather unexplored area (apart from a few publications so far). So the focus of the discussion has been on static constraints.

Linear functions play a key role in the area of rigorous runtime analysis of evolutionary computation. In [1], the runtime of evolutionary algorithms for linear functions with a linear constraint have been examined.

$$\max f(x) = \sum_{i=1}^{n} w_i x_i \text{ subject to } g(x) = \textstyle\sum_{i=1}^{n} g_i x_i \leq B.$$

In the general case, this is equivalent to the classical knapsack problem. However, it has been pointed out that the (1+1) EA which optimizes unconstrained linear functions in time $\Theta(n \log n)$ requires 2-bit flips even if the constrained is uniform, i.e. $g_i = 1$, $1 \leq i \leq n$. Understanding how even this simple case and provide tight upper and lower bounds would be of interest.

The participants mentioned different constrained handling mechanisms that could be examined such as a weightening of the given objective function f and the constrained g by considering the fitness given by

$$f(x) - \alpha \cdot \max\{0, g(x) - B\}.$$

An important question would be how to choose α (possibly adapting it over the run of the algorithm). Furthermore, multi-objective formulations (taking the constrained as an additional objective) may be examined.

Another topic discussed has been different black box models. For example, is the EA able to get information on the amount of constrained violation for an infeasible solutions or does it only get the information that a solution is infeasible. It's worthwhile comparing such different black box settings for prominent examples and analyze how evolutionary algorithms perform in the different settings.

Later there has been a discussion on dynamic changes to the constraints. One important question is what is the reoptimization time of evolutionary algorithms, i.e. the time to recompute a good or optimal solution after a change to the constraints has happened. Other questions involved the benefit of a population to cater for changes or reoptimize quicker. Furthermore, constrained handling mechanisms in relation to dynamic changes could be examined.

The breakout session has shown that there are a lot of open questions and interesting research directions for understanding how and why evolutionary algorithms can deal with constrained optimization problems.

References

1 T. Friedrich, T. Kötzing, G. Lagodzinski, F. Neumann, M. Schirneck: Analysis of the (1+1) EA on Subclasses of Linear Functions under Uniform and Linear Constraints. FOGA 2017: 45-54

4.6 Breakout Session: Diversity

Lothar Thiele (ETH Zürich, CH)

The breakout session was devoted to the concept of diversity. At the beginning, we discussed various areas in randomized search algorithms where the concept of diversity plays a major role. In particular, we agreed on the following classification:

Genotype: Diversity is a fundamental issue when discussing variation operators like crossover. In addition, it is used when arguing about adaptation and control, as well discussing about diversity in genotypes in order to deal with multi-modal optimization problems.

Phenotype: A novel and interesting optimization objective in many practical applications is the diversity of solutions, i.e. the goal to determine not only one solution but a set with different characteristics.

Objective Function: In multi-objective optimization, one major issue is to achieve diversity in the objective space. In particular, many methods are available to achieve a diverse set of solutions on or close to the Pareto Front.

Search Behavior: In complex search spaces, it is important to achieve a diverse search behavior, i.e. by using meta-search approaches. Moreover, other aspects of search behavior related to diversity are exploration vs. exploitation, initial solutions to randomized search algorithms, as well as examples for machine learning.

After discussing in some detail the above mention classification, possible definitions of diversity have been discussed that are suitable in several of the above instances:

- How well does a set cover a universe?
- Using a suitable information-theoretic measure, like Kolmogorov complexity.
- Following the concept of diversity as used in biology: Given is a completely connected undirected graph with edge weights. Adding a duplicate does not change the diversity, adding a distinct node increases the diversity, and increasing one of the edge weights also increases it.

Some of the interesting findings that should be considered further are the link to discrepancy theory and the associated concept of dispersion, and the fact that the desire of diversity is often a sign of uncertainty. The need for a general definition of diversity was questioned as (a) ad-hoc definitions worked well so far, and (b) diversity can also be controlled implicitly (selection pressure, mutation rate).

4.7 Breakout Session: COST Action CA15140

Christine Zarges (Aberystwyth University, GB)

Participants: Maxim Buzdalov, Arina Buzdalova, Francisco Chicano, Carola Doerr, Anton V. Eremeev, Carlos M. Fonseca, Thomas Jansen, William B. Langdon, Pietro S. Oliveto, Ofer M. Shir, Christine Zarges

The main aim of this breakout session was to discuss potential future collaborations within COST Action CA15140 "Improving Applicability of Nature-Inspired Optimisation by Joining Theory and Practice (ImAppNIO)".

4.7.1 Overview of the Action

COST is a European framework that provides funding for networking activities with an emphasis on Early Career Investigators (researchers with less than 8 years between the date of the PhD/doctorate and the date of involvement in the COST Action), inclusiveness and widening participation. Since most of the participants were not involved in the COST Action, yet, the first part of the meeting mainly aimed at introducing different activities and upcoming events and how new people can join.

4.7.1.1 Working Groups

Christine Zarges introduced the main aim of COST Action CA15140, which is chaired by Thomas Jansen: Build a platform where theoreticians and practitioners in nature-inspired optimisation can meet and exchange insights, ideas and needs. To achieve this, activities are planned along four different working groups:

- Working Group 1 (Theory-Driven Applications), led by Tobias Friedrich, is tasked with the development of novel theory-driven practical paradigms, thus pushing from theory to practice. Starting point are existing theoretical results and insights and the task is to use those to create practical guidelines, tangible advice and help for the application of nature-inspired search and optimisation heuristics.
- Working Group 2 (Practice-Driven Theory), led by Christine Zarges and Bosko Blagojevic, is tasked with the development of novel practice-driven theoretical frameworks and paradigms, thus pushing from practice to theory. Starting point are needs and unanswered questions as they arise in applications and the task is to create theoretical perspectives and novel results that directly address these needs.
- Working Group 3 (Benchmarks), led by Günther Raidl and Borys Wrobel, is tasked with the development of useful benchmarks for nature-inspired search and optimisation heuristics with a strong focus on discrete search spaces and discrete optimisation problems making sure that the developed benchmarks are relevant from a practical perspective and accessible from a theoretical perspective.
- Working Group 4 (Software), led by Carlos Fonseca and Florin Pop, is tasked with support for software development with a focus on the development of useful rules for the development of software and the adaptation of nature-inspired search and optimisation heuristics that are based on and guided by theoretical insights in their functioning.

More detailed information and current developments can be found on the COST action website: http://imappnio.dcs.aber.ac.uk

4.7.1.2 Training School

Carola Doerr introduced the upcoming COST training school, which will be centred around bridging the gap between theory and practice and making nature-inspired search and optimisation heuristics more applicable. It will take place from 18 to 24 October 2017 in Paris, France, right before the Biennial International Conference on Artificial Evolution (EA 2017), https://ea2017.inria.fr. Participation will be free and limited funding will be available for trainees from participating COST countries (see below). It is expected to be sufficient to pay for accommodation, subsistence and a significant contribution towards the cost of travel. For more details on the application process contact the action chair.

4.7.1.3 Short-Term Scientific Missions

Carola Doerr and Christine Zarges introduced the tool of Short-Term Scientific Missions (STSMs), exchange visits between researchers from two different countries involved in a COST Action. Applications for STSMs are invited at any time and need to be made via the e-COST system. More information about the process can be found on the COST action website: http://imappnio.dcs.aber.ac.uk/stsms

4.7.1.4 How to Participate?

The best first point of contact is the action chair, Thomas Jansen. To join a working group an additional email to the working group leader can be useful.

In a nutshell, COST distinguishes between three different types of countries: Member states (http://www.cost.eu/about_cost/cost_member_states), COST Near Neighbour Countries (http://www.cost.eu/about_cost/strategy/international_cooperation/nnc) and COST International Partner Countries (any other country). Any researcher affiliated with an institution in a member state already participating in the action is eligible for all activities. Researchers from non-participating member states need to contact the action chair and their national coordinator (http://www.cost.eu/about_cost/who/(type)/3) first to discuss how the state can join the action. The current list of participants can be found here: http://www.cost.eu/COST_Actions/ca/CA15140?parties. Institutions from COST Near Neighbour and International Partner Countries can join on a case by case basis and should discuss this with the action chair.

4.7.2 Discussion of Future Directions

In the remainder of the meeting, ideas for the training school including potential speakers were discussed in more detail. It was suggested to have ThRaSH-like talks towards the end of each day to make the school more attractive for senior researchers in the field who are not directly involved in the training.

A second discussion was centred around more concrete research ideas and potential routes to make working groups more effective. As a starting point participants shared their personal experiences including concrete collaborations and benchmarks. It was argued that issues such as language barriers or data protection could be overcome by first concentrating on problem modelling and presenting benchmarks as a black-box.

It was also suggested that the perceived gap between theory and practice is not symmetric as theory cannot expect to make significant progress in the available timeframe. Thus, an idea would be to ask practitioners to pose very simple questions to obtain a useful starting point and gain clarity about open questions. Here, it might be more promising to concentrate on collaborations with partners who are interested in developing methods to solve a specific *kind* of problem, rather than solutions for a very specific problem. However, in any case getting different researchers interested in theory and practice and collaborating on each others problems is a good starting point. A success story that solves a concrete problem and poses new questions would be desirable.

Finally, Carlos Fonseca gave a summary of the panel discussion at the COST industry workshop in Copenhagen earlier this year discussing requirements for evolutionary algorithms (e. g., speed, scalability, good default parameterisation), obstacles to adoption of evolutionary algorithms (e. g., lack of integration, trust and education) and aspects in favour of evolutionary algorithms (e. g., cost effectiveness, availability). It is hoped that wide collaboration in the COST Action can help to address some of these obstacles and generally improve the applicability of such methods.

5 Seminar Schedule

Monday, May 8

9.00 - 10.15	Welcome
	Introduction of Participants (40 participants, each 1–2 minutes)
10.15 - 10.45	Coffee Break
10.45 - 12.00	*Matt Hoffman*: Introductory Talk on Bayesian Optimization
12.15 - 14.00	Lunch Break
14.00 - 15.30	*Carsten Witt*: Recent Advances in Runtime Analysis of Estimation-of-Distribution Algorithms
	Olivier Teytaud: Randomized one-shot optimization
	J*on Rowe*: Linear multi-objective drift analysis
15.30 - 16.00	Coffee Break
16.00 - 18.00	*Benjamin Doerr*: Fast Genetic Algorithms
	Johannes Lengler: Noise Models for Comparison-Based EAs
	Bill Langdon: The fitness landscape of genetic improvement
	Thomas Jansen: COST Action

Tuesday, May 9

9.00 - 10.15	*Carlos M. Fonseca*: Mathematical Models of Artificial Genetic Representations with Neutrality
	Planning of the Breakout Sessions
10.15 - 10.45	Coffee Break
10.45 - 12.00	*Dirk Arnold*: Evolutionary optimization with constraints
	Dimo Brockhoff: Towards a Constrained Test Suite for COCO
	Tobias Glasmachers: Global Convergence of the (1+1)-ES
12.15 - 14.00	Lunch Break
14.00 - 15.30	Breakout Session I
	◦ Theory of evolutionary algorithms for problems with (dynamic) constraints
	◦ Drift theorems for continuous optimization - what we have vs. what we would need
	◦ Preference-based selection in evolutionary multiobjective optimization
15.30 - 16.00	Coffee Break
16.00 - 18.00	*Per Kristian Lehre*: Self-adaptation
	*Maxim Buzdalo*v: How to Exploit Your Fitness-Distance Correlation: Runtime Analysis of the (1+(?,?))-GA on Random Satisfiable 3-CNF Formulas
	Christian Giessen: Monotone Functions on Bitstrings - Some Structural Notes

Wednesday, May 10

9.00 - 10.15	*Nikolaus Hansen*: How to Guarantee Positive Definiteness in Active CMA-ES
	Hans-Georg Beyer: Towards a Theory of CMA-ES: But first Simplify your CMA-ES
10.15 - 10.45	Coffee Break
10.45 - 12.00	*Anton Eremeev*: Optimal Recombination for the TSP: Theory and Experiments
	Frank Neumann: Features, Diversity, Random Walks and Images
12.15	Lunch
13.30 - 15.30	Social Activity: Walk and Talk
15.30 - 16.30	Coffee Break
16.30 - 18.00	Breakout Session II
	◦ COST: Theory vs. Practice
	◦ Discrete estimation of distribution algorithms

Thursday, May 11

9.00 - 10.00	*Ofer Shir*: Fundamentals of evolution strategies' statistical learning
	Anne Auger: Connecting stability of Markov chains and deterministic control models for analyzing randomized algorithms
10.00 - 10.30	Coffee Break
10.30 - 12.00	*Pietro Oliveto*: Standard Steady State Genetic Algorithms can Hillclimb Faster than Mutation-Only Evolutionary Algorithms
	Günter Rudolph: Theoretical Aspects of the Averaged Hausdorff Indicator in Biobjective Optimization
	Group Photo
12.15 - 14.00	Lunch Break
14.00 - 15.30	Breakout Session III
	◦ Diversity in randomized optimization
	◦ Information geometric optimization. Topics: how to interpret the full CMA-ES in IGO
15.30 - 16.00	Coffee Break
16.00 - 18.00	*Francisco Chicano*: On the Variable Interaction Graph in Gray-Box Optimization
	Jing Yang: $(1+\lambda)$ Evolutionary Algorithm with Self-adjusting Mutation Rate
	Summary and Discussion of Breakout Sessions

Friday, May 12

9.15 - 10.15	*Jonathan Shapiro*: Max-Min Thompson Sampling for the K-arm dueling bandit problem
	Yohei Akimoto: Optimal Step-size for Weighted Recombination Evolution Strategy
	Bill Langdon: Long-term convergence vs long-term experimental evolution with genetic programming trees
10.15 - 10.45	Coffee Break
10.45 - 12.00	*Martin Krejca*: Estimation of Distribution Algorithms
	Arina Buzdalova: Is it necessary to perform multi-objective optimization when doing multi-objectivization?
	Wrap-up
12.15	Lunch

Participants

- Youhei Akimoto
Shinshu University – Nagano, JP
- Dirk V. Arnold
Dalhousie University – Halifax, CA
- Anne Auger
INRIA RandOpf Team, FR
- Thomas Bäck
Leiden University, NL
- Hans-Georg Beyer
Fachhochschule Vorarlberg – Dornbirn, AT
- Dimo Brockhoff
INRIA Saclay – Île-de-France, FR
- Maxim Buzdalov
ITMO University – St. Petersburg, RU
- Arina Buzdalova
ITMO University – St. Petersburg, RU
- Francisco Chicano
University of Málaga, ES
- Benjamin Doerr
Ecole Polytechnique – Palaiseau, FR
- Carola Doerr
CNRS & UPMC, Paris, FR
- Anton V. Eremeev
Sobolev Institute of Mathematics – Novosibirsk
- Carlos M. Fonseca
University of Coimbra, PT
- Tobias Friedrich
Hasso-Plattner-Institut – Potsdam, DE
- Christian Gießen
Technical University of Denmark – Lyngby, DK
- Tobias Glasmachers
Ruhr-Universität Bochum, DE
- Nikolaus Hansen
INRIA Saclay – Île-de-France, FR
- Matthew W. Hoffman
Google DeepMind – London, GB
- Christian Igel
University of Copenhagen, DK
- Thomas Jansen
Aberystwyth University, GB
- Martin S. Krejca
Hasso-Plattner-Institut – Potsdam, DE
- William B. Langdon
University College London, GB
- Per Kristian Lehre
University of Birmingham, GB
- Johannes Lengler
ETH Zürich, CH
- Frank Neumann
University of Adelaide, AU
- Pietro S. Oliveto
University of Sheffield, GB
- Adam Prugel-Bennett
University of Southampton, GB
- Jonathan E. Rowe
University of Birmingham, GB
- Günter Rudolph
TU Dortmund, DE
- Jonathan L. Shapiro
University of Manchester, GB
- Ofer M. Shir
Tel-Hai College – Upper Galilee, IL
- Dirk Sudholt
University of Sheffield, GB
- Andrew M. Sutton
Hasso-Plattner-Institut – Potsdam, DE
- Olivier Teytaud
Google Switzerland – Zürich, CH
- Lothar Thiele
ETH Zürich, CH
- Heike Trautmann
Universität Münster, DE
- Carsten Witt
Technical University of Denmark – Lyngby, DK
- Jing Yang
Ecole Polytechnique – Palaiseau, FR
- Xin Yao
University of Birmingham, GB
- Christine Zarges
Aberystwyth University, GB

Report from Dagstuhl Seminar 17192

Human-Like Neural-Symbolic Computing

Edited by

Tarek R. Besold[1], Artur d'Avila Garcez[2], and Luis C. Lamb[3]

1 **Universität Bremen, DE,** tbesold@uni-bremen.de
2 **City, University of London, GB,** aag@soi.city.ac.uk
3 **Federal University of Rio Grande do Sul, BR,** luislamb@acm.org

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17192 "Human-Like Neural-Symbolic Computing", held from May 7th to 12th, 2017. The underlying idea of Human-Like Computing is to incorporate into Computer Science aspects of how humans learn, reason and compute. Whilst recognising the relevant scientific trends in big data and deep learning, capable of achieving state-of-the-art performance in speech recognition and computer vision tasks, limited progress has been made towards understanding the principles underlying language and vision understanding. Under the assumption that neural-symbolic computing – the study of logic and connectionism as well statistical approaches – can offer new insight into this problem, the seminar brought together computer scientists, but also specialists on artificial intelligence, cognitive science, machine learning, knowledge representation and reasoning, computer vision, neural computation, and natural language processing. The seminar consisted of contributed and invited talks, breakout and joint group discussion sessions, and a hackathon. It was built upon previous seminars and workshops on the integration of computational learning and symbolic reasoning, such as the Neural-Symbolic Learning and Reasoning (NeSy) workshop series, and the previous Dagstuhl Seminar 14381: Neural-Symbolic Learning and Reasoning.

Seminar May 7–12, 2017 – http://www.dagstuhl.de/17192

1998 ACM Subject Classification F.1.1 Models of Computation: Self-modifying machines, I.2 Artificial Intelligence, I.2.10 Vision and Scene Understanding, I.2.4 Knowledge Representation Formalisms and Methods, I.2.6 Learning, I.2.6 Learning: Connectionism and neural nets, I.2.7 Natural Language Processing

Keywords and phrases Deep Learning, Human-like computing, Multimodal learning, Natural language processing, Neural-symbolic integration

Digital Object Identifier 10.4230/DagRep.7.5.56

1 Executive Summary

Luis C. Lamb
Tarek R. Besold
Artur d'Avila Garcez

The underlying idea of Human-Like Computing is to incorporate into Computer Science aspects of how humans learn, reason and compute. Recognising the relevance of the scientific trends in big data, data science methods and techniques have achieved industrial relevance in a number of areas, from retail to health, by obtaining insight from large data collections. Notably, neural networks have been successful and efficient at large-scale language modelling, speech recognition, image, video and sensor data analysis [3, 12, 15]. Human beings, on the other hand, are excellent at learning from very few data examples, capable of articulating explanations and resolving inconsistencies through reasoning and communication [7, 9, 16, 17].

Despite the recent impact of deep learning, limited progress has been made towards understanding the principles and mechanisms underlying language and vision understanding. Under this motivation, the seminar brought together not only computer scientists, but also specialists in artificial intelligence (AI), cognitive science, machine learning, knowledge representation and reasoning, computer vision, neural computation and natural language processing. In particular, the methodology of neural-symbolic computation [4, 7, 12], which can offer a principled interface between the relevant fields, especially symbolic AI and neural computation, was adopted in an attempt to offer a new perspective of reconciling large-scale modelling with human-level understanding, thus building a roadmap for principles and applications of Human-Like Neural-Symbolic Computing.

The techniques and methods of neural-symbolic computation have already been applied effectively to a number of areas, leading to developments in deep learning, data science and human-like computing [3]. For instance, neural-symbolic integration methods have been applied to temporal knowledge evolution in dynamic scenarios [6, 10, 14], action learning in video understanding [10], uncertainty learning and reasoning [1], argument learning in multiagent scenarios [7, 8], hardware and software verification and learning [2], ontology learning [13] and distributed temporal deep learning in general, with several applications in computer science [2, 10, 14].

Specifically, in this Dagstuhl Seminar we aimed at: (i) building better bridges between symbolic and sub-symbolic reasoning and learning, and between big data and human-like learning; (ii) comparative analyses and evaluations of the explanatory capacity of language modelling tools and techniques; (iii) design and applications of knowledge extraction methods and techniques towards life-long learning and transfer learning.

The seminar consisted of contributed and invited talks, breakout and joint group discussion sessions, and scientific hackathons. After each presentation or discussion session, open problems were identified and questions were raised. The area is clearly growing in importance, given recent advances in Artificial Intelligence and Machine Learning. In particular, the need for explainability in AI clearly poses relevant questions to learning methodologies, including deep learning. In summary, the main research directions identified by participants are:

- Explainable AI: The recent success of deep learning in vision and language processing, associated with the growing complexity of big data applications has led to the need for *explainable AI models*. In neural-symbolic computing, rule extraction, interpretability, comprehensibility leading to the development of integrated systems, are one of the principled alternatives to lead these efforts [5, 7], as discussed in the Explainability hackathon. Furthermore, the concept of modularity in multimodal learning in deep networks is crucial to the development of the field and can help achieve knowledge extraction (as identified in [5, 6]) which can result in the development of effective knowledge extraction methods towards explainable AI, as discussed in the deep learning with symbols hackathon.
- Hybrid Cognitive Architectures: The development of Cognitive Architectures capable of simulating and explaining aspects of human cognition also remains an important research endeavour. Some cognitive architectures typically consider symbolic representations, whereas others employ neural simulations. The integration of these models remains a challenge and there are benefits on integrating the accomplishments of both paradigms, as identified in the cognitive architectures hackathon.

- Statistical Relational Learning: Logic Tensor Networks (LTNs) [11] provides a model that integrates symbolic knowledge (encoded as first-order logic relations) and subsymbolic knowledge (represented as feature vectors). The LTNs enable the representation of relational knowledge infusion into deep networks, and knowledge completion and distilling through querying the networks. There remains a number of challenges in integrating, explaining and computing symbolic knowledge in deep networks. Both LTNs [11] and Connectionist Modal and Temporal Logics [6, 7, 14] offer effective alternatives towards these research challenges, as explored in the LTN hackathon.

The seminar builds upon previous seminars and workshops on the integration of computational learning and symbolic reasoning, such as the Neural-Symbolic Learning and Reasoning (NeSy) workshop series, and the previous Dagstuhl Seminar 14381: Neural-Symbolic Learning and Reasoning [5].

References

1 Tarek R. Besold, Artur S. d'Avila Garcez, Keith Stenning, Leendert W. N. van der Torre, Michiel van Lambalgen: Reasoning in Non-probabilistic Uncertainty: Logic Programming and Neural-Symbolic Computing as Examples. *Minds and Machines*, 27(1): 37-77, 2017.

2 Rafael V. Borges, Artur S. d'Avila Garcez, Luis C. Lamb. *Learning and Representing Temporal Knowledge in Recurrent Networks*. IEEE Trans. Neural Networks and Learning Systems 22(12):2409-2421, Dec. 2011. IEEE Press.

3 A. d'Avila Garcez, Tarek R. Besold, Luc De Raedt, Peter Földiak, Pascal Hitzler, Thomas Icard, Kai-Uwe Kühnberger, Luis C. Lamb, Risto Miikkulainen, Daniel L. Silver. Neural-Symbolic Learning and Reasoning: Contributions and Challenges. Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches. Stanford Univ., March 2015, pp. 18-21, AAAI press, 2015.

4 Artur S. d'Avila Garcez, Dov M. Gabbay, Krysia Broda. *Neural-Symbolic Learning System: Foundations and Applications*. Springer-Verlag, New York, Inc., USA, 2002.

5 A. S. d'Avila Garcez, Marco Gori, Pascal Hitzler, Luís C. Lamb: Neural-Symbolic Learning and Reasoning (Dagstuhl Seminar 14381). Dagstuhl Reports 4(9): 50-84, 2015.

6 Artur S. d'Avila Garcez and Luis C. Lamb. A Connectionist Computational Model for Epistemic and Temporal Reasoning. *Neural Computation*, 18(7):1711-1738, 2006.

7 Artur S. d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay. *Neural-Symbolic Cognitive Reasoning*. Springer Publishing Company, 2009. ISBN: 9783540732457.

8 A. d'Avila Garcez, D.M. Gabbay and Luis C. Lamb: A neural cognitive model of argumentation with application to legal inference and decision making. Journal of Applied Logic, 12(2):109-127, 2014.

9 M. de Kamps and F. van de Velde. From neural dynamics to true combinatorial structures. Behavioral and Brain Sciences, 20, pp. 88-108, 2006.

10 H.L. de Penning, Artur S. d'Avila Garcez, Luis C. Lamb, John-Jules Ch. Meyer: *A Neural-Symbolic Cognitive Agent for Online Learning and Reasoning*. Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence, IJCAI 2011. Pages 1653-1658, 2011.

11 Ivan Donadello, Luciano Serafini, and Artur d'Avila Garcez. *Logic Tensor Networks for Semantic Image Interpretation*. Proceedings of the 26th International Joint Conference on Artificial Intelligence, 2017.

12 Barbara Hammer, Pascal Hitzler (Eds.): Perspectives of Neural-Symbolic Integration. Studies in Computational Intelligence 77, Springer 2007.

13 P. Hitzler, S. Bader, and A. d'Avila Garcez. *Ontology learning as a use case for neural-symbolic integration*. In Proc. Workshop on Neural-Symbolic Learning and Reasoning, NeSy'05 at IJCAI-05.

14 Luis C. Lamb, R.V. Borges and A.S. d'Avila Garcez. *A connectionist cognitive model for temporal synchronisation and learning*. Proc. AAAI Conference on Artificial Intelligence AAAI-07, pages 827-832, AAAI Press, 2007.
15 Y. LeCun, Y. Bengio, G. Hinton. Deep Learning. *Nature*, 521 (7553): 436-444, 2015.
16 G.F. Marcus, S. Vijayan, S.B. Rao, and P.M. Vishton. Rule learning by seven-month-old infants. *Science*, 283(5398):77–80, 1999.
17 K. Stenning and M. van Lambalgen. *Human reasoning and Cognitive Science*. MIT Press, Cambridge, MA, 2008.
18 Alan M. Turing. Computing Machinery and Intelligence, *Mind*, LIX (236): 433–460, 1950.

2 Table of Contents

Executive Summary
Luis C. Lamb, Tarek R. Besold, and Artur d'Avila Garcez 56

Overview of Talks

Conceptual Spaces: A Bridge Between Neural and Symbolic Representations?
Lucas Bechberger 62

Comprehensible ILP
Tarek R. Besold 62

Using evidence accumulation to bridge the gap between neural networks and symbolic cognitive control
Jelmer Borst and Terrence Stewart 63

Neural-Symbolic Systems for Human-Like Computing
Artur d'Avila Garcez 63

Dynamics for the Neural Blackboard Architecture
Marc de Kamps 64

On Explainability in Machine Learning
Derek Doran 66

Tackling Commonsense Reasoning Benchmarks
Ulrich Furbach 66

What do Neural Networks need in order to generalize?
Raquel Garrido Alhama and Willem Zuidema 66

Understanding Neural Networks through Background Knowledge
Pascal Hitzler, Derek Doran, Maryam Labaf, Md Kamruzzaman Sarker, Michael Raymer, and Ning Xie 67

Processing Hierarchical Structure with RNNs
Dieuwke Hupkes 68

How do people build and comprehend ontologies?
Caroline Jay 68

Human-like software engineering: bridging the human-machine translation gap
Caroline Jay 69

(Human-Like) Anecdotes with IBM's Watson/Bluemix Services
Kai-Uwe Kühnberger 70

From Turing to Deep Learning: Explaining AI through neurons and symbols
Luis C. Lamb 71

Interacting Conceptual Spaces
Martha Lewis 72

Deep Learning with Symbols
Daniel L. Silver 73

Grounded Language Processing and Learning
Michael Spranger 73

Are 'logically naive' people logically naive?
Keith Stenning 74

Logic Programming as a framework for understanding human computation
Keith Stenning 75

Structured Computer Organization of the Human Mind
Niels A. Taatgen 75

Embodied neuro-symbolic computation
Serge Thill 76

Overview of neuro-symbolic processing in Neural Blackboard Architectures
Frank Van der Velde 76

Neural Computing with Signals and Symbols for Music Analysis
Tillman Weyde 77

Working groups

Logic Tensor Networks Hackathon
Lucas Bechberger, Artur d'Avila Garcez, Raquel Garrido Alhama, Marco Gori, Luciano Serafini, Michael Spranger, and Tillman Weyde 78

Explainability Hackathon
Tarek R. Besold, Derek Doran, Pascal Hitzler, Caroline Jay, Oliver Kutz, and Sarah Schulz 79

Deep Learning with Symbols Hackathon
Daniel L. Silver, James Christopher Davidson, Dieuwke Hupkes, Isaac Noble, and Katja Seeliger 80

Nengo/ACT-R/PRIMs/Cognitive Architectures Hackathon
Niels A. Taatgen, Jelmer Borst, Marc de Kamps, Martha Lewis, Serge Thill, and Frank Van der Velde 81

Open problems

Research Challenges
Luis C. Lamb, Tarek R. Besold, and Artur d'Avila Garcez 81

Participants 83

3 Overview of Talks

3.1 Conceptual Spaces: A Bridge Between Neural and Symbolic Representations?

Lucas Bechberger (Universität Osnabrück, DE)

The cognitive framework of conceptual spaces [1] attempts to bridge the gap between symbolic and subsymbolic AI by proposing an intermediate conceptual layer based on geometric representations. A conceptual space is a high-dimensional space spanned by a number of quality dimensions representing interpretable features. Convex regions in this space correspond to concepts. Abstract symbols can be grounded by linking them to concepts in a conceptual space whose dimensions are based on subsymbolic representations.

The framework of conceptual spaces has been highly influential in the last 15 years within cognitive science and cognitive linguistics. It has also sparked considerable research in various subfields of artificial intelligence, ranging from robotics and computer vision over the semantic web and ontology integration to plausible reasoning.

Although this framework provides means for connecting concepts from the symbolic layer to numeric information from the subsymbolic layer, it does not yet provide an automated way to do so: In practical applications, both the mapping from the subsymbolic to the conceptual layer (i.e., how to map observations to points) and the mapping from the conceptual layer to the symbolic layer (i.e., how to map regions to symbols) need to be handcrafted by a human expert.

After introducing the conceptual spaces framework in more detail, I argue that we can use machine learning in order to learn these mappings: I propose to use representation learning (namely the InfoGAN framework [2]) for learning the dimensions of a conceptual space from unlabeled data. Moreover, I suggest to use an incremental clustering algorithm to discover meaningful regions in a conceptual space that can then give rise to abstract symbols.

References

1 Peter Gärdenfors. *Conceptual Spaces: The Geometry of Thought*. MIT Press, 2000
2 Xi Chen, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel, *InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets*. Advances in Neural Information Processing Systems, 2016.

3.2 Comprehensible ILP

Tarek R. Besold (Universität Bremen, DE)

Joint work of Ute Schmid, Christina Zeller, Stephen Muggleton, Alireza Tamaddoni-Nezhad, Tarek R. Besold
Main reference Schmid U., Zeller C., Besold T., Tamaddoni-Nezhad A., Muggleton S. (2017) How Does Predicate Invention Affect Human Comprehensibility?. In: Cussens J., Russo A. (eds) Inductive Logic Programming. ILP 2016. Lecture Notes in Computer Science, vol 10326. Springer, Cham
URL http://dx.doi.org/10.1007/978-3-319-63342-8_5

During the 1980s Michie defined Machine Learning in terms of two orthogonal axes of performance: predictive accuracy and comprehensibility of generated hypotheses. Since predictive accuracy was readily measurable and comprehensibility not so, later definitions in the 1990s, such as that of Mitchell, tended to use a one-dimensional approach to Machine Learning based solely on predictive accuracy, ultimately favoring statistical over symbolic Machine Learning approaches. In this talk I want to discuss a definition of comprehensibility of hypotheses which can be estimated using human participant trials.

To illustrate the proposed approach, we will have a look at two recent experiments testing human comprehensibility of ILP-learned logic programs. The first experiment tested human comprehensibility with and without predicate invention, with results indicating that comprehensibility is affected not only by the complexity of the presented program but also by the existence of anonymous predicate symbols. The second experiment showed that a state-of-the-art ILP system can support humans effectively beyond their own capabilities in a relational concept learning task.

3.3 Using evidence accumulation to bridge the gap between neural networks and symbolic cognitive control

Jelmer Borst (University of Groningen, NL) and Terrence Stewart

I will discuss how we used large-scale spiking neural networks to simulate associative recognition. Associative recognition is the important ability to learn that two items co-occur. Although detailed symbolic models exist that account for behavior, fMRI, and EEG data, it remains unclear how associative recognition is performed at the neural level. To investigate this, we used the Neural Engineering Framework to simulate associative recognition with spiking neural networks that can process symbols and coordinate cognition through the basal ganglia [2]. Because the resulting neural network model is very complex (> 500,000 neurons) we use magnetoencephalographic (MEG) data to constrain the model [1]. The model matches data in occipital, temporal, prefrontal, and motor cortices, and shows how the associative recognition process could be implemented in the human brain.

References

1 Jelmer P. Borst, Avniel S. Ghuman, & John R. Anderson *Tracking cognitive processing stages with MEG: A spatio-temporal model of associative recognition in the brain*. NeuroImage, 141:416–430, 2016.

2 Chris Eliasmith, Terrence C. Stewart, Xuan Choo, Trevor Bekolay, Travis DeWolf, Yichuan Tang, & Daniel Rasmussen *A Large-Scale Model of the Functioning Brain*. Science, 338:1202–1205, 2012.

3.4 Neural-Symbolic Systems for Human-Like Computing

Artur d'Avila Garcez (City, University of London, GB)

An increasing number of researchers now accept that the recent success of deep learning needs to be combined with an understanding of its reach, scope and limitations so that the area can continue to move forward towards Artificial General Intelligence (AGI). Neural-symbolic computing has been making foundational contributions in this direction for a long time now [1, 3], as it is concerned with the interplay between learning and reasoning, and the combination of rich symbolic AI formalisms with efficient neural computation frameworks. The applications of AGI are vast as it permits humans to improve their own performance at various tasks from interacting with (explainable) learning systems. AGI also permits systems to learn from other systems as part of what became known as transfer learning. Finally, such general learning systems will need to become verifiable against certain properties such as safety and trust as they become commonplace in our daily lives [2].

In "Neural-Symbolic Systems for Human-Like Computing", I presented an overview of the neural-symbolic methodology, including knowledge insertion (aka infusion), learning, reasoning and knowledge extraction (aka distilling) from neural networks. This methodology enables the integration of symbolic and sub-symbolic AI in a principled way. I then exemplified the neural-symbolic cycle by applying it to software model verification and adaptation [5], and to the use of neural networks for the run-time monitoring of software system properties specified in linear temporal logic [4]. I concluded by discussing ways in which neural-symbolic systems may contribute to the research on human-like computing, notably in the areas of representation change, bridging the gap between high and low-level learning and reasoning, memory and forgetting, comprehensibility and explanation, and learning from only a few examples.

References

1. Garcez, Artur S. d'Avila, Gabbay, Dov M. and Broda, Krysia B., *Neural-Symbolic Learning System: Foundations and Applications*. Springer-Verlag, New York, Inc., USA, 2002. ISBN: 1852335122.
2. Garcez, Artur S. d'Avila and Lamb, Luis C., *Reasoning About Time and Knowledge in Neural-symbolic Learning Systems*. In Proc. 16th International Conference on Neural Information Processing Systems NIPS'03, MIT Press, Cambridge, MA, USA, 2003.
3. Garcez, Artur S. d'Avila, Lamb, Luis C. and Gabbay, Dov M., *Neural-Symbolic Cognitive Reasoning*. Springer Publishing Company, 2008. ISBN: 9783540732457.
4. Alan Perotti, Artur S. d'Avila Garcez and Guido Boella, *Neural-symbolic monitoring and adaptation*. In Proc. International Joint Conference on Neural Networks, IJCNN'15, Killarney, Ireland, 2015.
5. Borges, Rafael V., d'Avila Garcez, Artur, Lamb, Luis C., *Learning and Representing Temporal Knowledge in Recurrent Networks*. IEEE Trans. Neural Networks and Learning Systems 22(12):2409–2421, Dec. 2011. IEEE Press.

3.5 Dynamics for the Neural Blackboard Architecture

Marc de Kamps (University of Leeds, GB)

Humans surpass computers when it comes to understanding. Computers excel in object recognition or sentence parsing, but when it comes to extracting meaning, humans do a lot better. We partly understand why this is so: neural representations are grounded [4]: neural activity derives meaning from where it occurs relative to the sensory pathways or other brain areas. This is a major advantage of computer representations consisting of strings of binary numbers, which have no intrinsic meaning. The network knowledge representation of the brain makes it relatively straightforward to understand how associations occur and how attractor networks can implement memory. This implementation raises fundamental questions about how compositional representations are represented, however. It is not easy to understand how the same concept can occur twice in a different context within the same scene or sentence [6]. It is not easy to understand how abstract relationships are realized by the brain and how they can be applied to concrete instances, when an agent must resolve its constituents in order to extract meaning To address this problem, we proposed the Neural Blackboard Architecture (NBA), a biologically plausible cognitive architecture that aims to explain these capabilities in terms of neural dynamics [5]. We have reached the point where we can make predictions about sequences of neural activity in language processing, and validate them against experimental results. Preliminary results of this work have been reported in abstract form [2].

For such a program, realistic models of neural dynamics at the population level are essential. I will report recent progress in population density techniques (e.g. [1, 3]), developed as part of the Human Brain Project, to show how large groups of spiking neurons can be modeled efficiently at the population level, without using thousands of model neuron instances. I will also show preliminary modeling results of this technique applied to the NBA. When four word sentences are compared with a baseline consisting of four words not constituting a sentence, there is a marked difference in the neural signal. Future work will make the comparison to experimental results. A big advantage of the technique presented here is that neural models can be exchanged easily without affecting the structure of the network, which makes it possible to vary biological assumptions in the model (e.g. the presence of adaptation). The population density technique implementation is available as Open Source software on http://miind.sf.net

Acknowledgment: The population modeling work received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 720270 (Human Brain Project).

References

1 Y.M. Lai and M. de Kamps. Population density equations for stochastic processes with memory kernels, Phys. Rev. E 95, 062125, 2017.
2 Martin Perez-Guevara, Marc De Kamps, Christophe Pallier. Blackboard Architecture simulation with simple biological networks captures the behavior of diverse neuroimaging measurements during language processing, in NIPS Workshop Representation Learning in Artificial and Biological Neural Networks, 2016.
3 de Kamps M, A Generic Approach to Solving Jump Diffusion Equations with Applications to Neural Populations, https://arxiv.org/abs/1309.1654, 2013.
4 F. van der Velde, M. de Kamps. Compositional connectionist structures based on in situ grounded representations. Connect Sci, 23 (2), pp. 97-107, 2011.
5 M. de Kamps, F. van de Velde. From neural dynamics to true combinatorial structures. Behavioral and Brain Sciences, 20 , pp. 88-108, 2006.
6 R. Jackendoff. Foundations of Language: Brain, Meaning, Grammar, Evolution, Oxford University Press, 2003.

3.6 On Explainability in Machine Learning

Derek Doran (Wright State University – Dayton, US)

Joint work of Ning Xie, Md Kamruzzaman Sarker, Eric Nicols, Pascal Hitzler, Michael Raymer, Derek Doran

Enabling explainability in statistical machine learning is a fast emerging topic. The development is rooted in the idea that statistical ML models are widely perceived (perhaps unfairly) as "black box", where input/output mappings are performed by mathematical operations that are not easily interpretable or explainable. But (many) statisticians and ML researchers disagree: one can argue that regression models are interpretable, and any model with linear decision boundaries for classification imply a set of decision making "rules". This presentation introduces a hierarchy of explainability that statistical ML algorithms may be classified into: interpretable, explainable, and reasonable, and includes recent examples of each. It then introduces a new approach for enabling explanations, and potentially higher level reasoning, for neural networks. The approach involves: 1) localized representations of concepts through regularization; 2) applying an inference engine that explains these localized activations by discriminating between inputs activating the localization and those that don't.
Acknowledgements: This work is supported by the Ohio Federal Research Network.

3.7 Tackling Commonsense Reasoning Benchmarks

Ulrich Furbach (Universität Koblenz-Landau, DE)

Joint work of Ulrich Furbach, Claudia Schon
Main reference U. Furbach, C. Schon, "Commonsense Reasoning Meets Theorem Proving", in Proc. of the 14th German Conference on Multiagent System Technologies (MATES 2016), LNCS, Vol. 9872, pp. 3–17, Springer, 2016.
URL https://doi.org/10.1007/978-3-319-45889-2_1

The area of commonsense reasoning aims at the creation of systems able to simulate the human way of rational thinking. This talk describes the use of automated reasoning methods for tackling commonsense reasoning benchmarks. For this we use a benchmark suite introduced in literature. Our goal is to use general purpose background knowledge without domain specific hand coding of axioms, such that the approach and the result can be used as well for other domains in mathematics and science. Furthermore, we discuss the modeling of normative statements in commonsense reasoning and in robot ethics.

3.8 What do Neural Networks need in order to generalize?

Raquel Garrido Alhama (University of Amsterdam, NL) and Willem Zuidema

In an influential paper, reporting on a combination of artificial language learning experiments with babies, computational simulations and philosophical arguments, [1] claimed that connectionist models cannot account for human success at learning tasks that involved generalization of abstract knowledge such as grammatical rules. This claim triggered a heated debate, centered mostly around variants of the Simple Recurrent Network model [2]. In this paper, we revisit this unresolved debate and analyze the underlying issues from a different perspective. We argue that, in order to simulate human-like learning of grammatical rules, a neural network model should not be used as a tabula rasa, but rather, the initial wiring of the neural connections and the experience acquired prior to the actual task should be incorporated into the model. We present two methods that aim to provide such initial state: a manipulation of the initial connections of the network in a cognitively plausible manner (concretely, by implementing a "delay-line" memory), and a pre-training algorithm that incrementally challenges the network with novel stimuli. We implement such techniques in an Echo State Network [3], and we show that only when combining both techniques the ESN is able to succeed at the grammar discrimination task suggested by Marcus et al.

References

1 G. F. Marcus, S. Vijayan, S.B. Rao, and P.M. Vishton. Rule learning by seven-month-old infants. *Science*, 283(5398):77–80, 1999.

2 Jeffrey L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.

3 Herbert Jaeger. *The "echo state" approach to analysing and training recurrent neural networks*. Technical report, German National Research Center for Information Technology, 2001.

3.9 Understanding Neural Networks through Background Knowledge

Pascal Hitzler (Wright State University – Dayton, US), Derek Doran (Wright State University – Dayton, US), Maryam Labaf, Md Kamruzzaman Sarker, Michael Raymer, and Ning Xie

The ever increasing prevalence of publicly available structured data on the World Wide Web enables new applications in a variety of domains. In this presentation, we provide a conceptual approach that leverages such data in order to explain the input-output behavior of trained artificial neural networks. We apply existing Semantic Web technologies in order to provide an experimental proof of concept. The presentation starts by investigating the case of propositional rule extraction as a base case, carrying past results over to incorporate background knowledge. In the second part, we incorporate knowledge graphs and ontologies and show how the DL-Learner symbolic machine learning system can be used to generate explanations which take such background knowledge into account.

Acknowledgements. This work was supported by the Ohio Federal Research Network project *Human-Centered Big Data*.

References

1 Maryam Labaf and Pascal Hitzler, Propositional Rule Extraction from Neural Networks under Background Knowledge. In: Proceedings of the Twelfth International Workshop on Neural-Symbolic Learning and Reasoning, London, UK, July 2017. To appear.

2 Md Kamruzzaman Sarker and Ning Xie and Derek Doran and Michael Raymer and Pascal Hitzler, Explaining Trained Neural Networks with Semantic Web Technologies: First Steps. In: Proceedings of the Twelfth International Workshop on Neural-Symbolic Learning and Reasoning, London, UK, July 2017. To appear.

3.10 Processing Hierarchical Structure with RNNs

Dieuwke Hupkes (University of Amsterdam, NL)

Joint work of Willem Zuidema, Sara Veldhoen, Dieuwke Hupkes

We investigate how recurrent neural networks can learn and process languages h hierarchical, compositional semantics. To this end, we define the simple but nontrivial artificial task of processing nested arithmetic expressions and study whether recurrent neural networks, which process these expressions incrementally, can learn to compute their meaning. We show that a neural network architecture with gating (the Gaed Recurrent Unit) performs well on this task: the network learns to predict the outcome of the arithmetic expressions with high accuracy, although performance deteriorates with increasing length. To analyse what strategy the recurrent network applies, visualisation techniques are not sufficient. Therefore, we develop an approach where we formulate and test hypothess on what strategies such networks might be following. For each hypothesis, we derive predictions about featurs of the hidden state representations at each time step, and train 'diagnostic' classifiers to test those predictions. Our results indicate that the networks follow a strategy similar to our hypothesised 'cumultative strategy'.

3.11 How do people build and comprehend ontologies?

Caroline Jay (University of Manchester, GB)

Joint work of Caroline Jay, Nicolas Matentzoglu, Markel Vigo, Robert Stevens
Main reference N. Matentzoglu, M. Vigo, C. Jay, R. Stevens, "Making Entailment Set Changes Explicit Improves the Understanding of Consequences of Ontology Authoring Actions", in Proc. of the 20th Int'l Conf. on Knowledge Engineering and Knowledge Management (EKAW 2016), LNCS, Vol. 10024, pp. 432–446, Springer, 2016.
URL https://doi.org/10.1007/978-3-319-49004-5_28

Ontologies have become an important feature of many fields, including science, commerce and medicine. They support the representation of a domain's knowledge for use in a range of applications. Ontologies describe classes of objects, and the relationships between them. When in logical form, these descriptions can be used to draw conclusions about the knowledge represented; this information can be used both for developing the ontology itself, and in an application. An ontology can become very large and highly complex, however, and a single change can have many unanticipated effects (both desirable and undesirable) that are difficult for a human engineer to comprehend [2]. Accuracy is important: if the representation is incorrect, the value of any knowledge resulting from reasoning over it is undermined.

At the University of Manchester, via the EPSRC-funded project 'WhatIf: Answering "What if. . . " questions for Ontology Authoring' (EP/J014176/1), we have been exploring how people build ontologies, and comprehend the information within them. This work shows that people use the class hierarchy (within the Protege tool) as an external memory, spending almost half their time looking or interacting with it when working on an ontology. They also use the reasoner and the inferred class hierarchy to check the consequences of their modelling actions [3].

Understanding how a particular action has affected an ontology is extremely difficult, as there may be multiple, unanticipated changes that have implications for a complex series of axioms. To ameliorate the situation, we designed and tested the Inference Inspector, a tool that shows the most relevant consequences of modelling actions to the ontology developer. Making entailment set changes explicit significantly improves the speed and accuracy of ontology development, as well as improving author comprehension of the effects of his or her changes [1].

References

1 Nicolas Matentzoglu, Markel Vigo, Caroline Jay and Robert Stevens. *Making Entailment Set Changes Explicit Improves the Understanding of Consequences of Ontology Authoring Actions.* Knowledge Engineering and Knowledge Management – 20th International Conference EKAW 2016, Bologna, Italy, November 19–23, 2016.

2 Markel Vigo, Sam Bail, Caroline Jay and Robert Stevens. *Overcoming the pitfalls of ontology authoring: Strategies and implications for tool design.* International Journal of Human-Computer Studies, 2014.

3 Markel Vigo, Caroline Jay and Robert Stevens. *Constructing Conceptual Knowledge Artefacts: Activity Patterns in the Ontology Authoring Process.* CHI, 2015.

3.12 Human-like software engineering: bridging the human-machine translation gap

Caroline Jay (University of Manchester, GB)

Joint work of Caroline Jay, Robert Haines, Robert Stevens
Main reference C. Jay, R. Haines, M. Vigo, N. Matentzoglu, R. Stevens, J. Boyle, A. Davies, C. Del Vescovo, N. Gruel, A. Le Blanc, D. Mawdsley, D. Mellor, E. Mikroyannidi, R. Rollins, A. Rowley, J. Vega Hernandez, "Identifying the challenges of code/theory translation: report from the Code/Theory 2017 workshop", Research Ideas and Outcomes 3: e13236, Pensoft Publishers, 2017.
URL http://dx.doi.org/10.3897/rio.3.e13236

Engineering software is a challenging endeavour. Development processes are incrementally improving, allowing us to construct increasingly complex artefacts, yet software continues to contain errors, or behave in unforeseen ways. This is partly to due to the 'unknown unknowns' introduced by a changing external environment, but it is also because algorithms often fail to work as expected: the formal representations underlying machine computation are frequently at odds with the heuristics used by the human brain. Whilst this is an issue across software development, it is particularly apparent in scientific software engineering, where the purpose of the code implementation is to represent, precisely, a scientific entity, process or system, and a given programming language is often an inadequate means of doing this [1]. The result is a *human-machine translation gap.*

Whilst empirical software engineering is a thriving field, using a theory of human cognition based on empirical observation to drive system development is not a standard approach; progress has occurred primarily through craft-based iteration, rather than rigorous empirical study. Nevertheless, observation of the programming process has demonstrated its potential to result in huge technological advances. A notable example of this is locality of reference, a principle uncovered when trying to ascertain how to page data in and out of memory, which has gone on to touch virtually every aspect of modern systems. The *theory of locality* describes how data relevant to the current context of a running program is grouped together locally in space and time. Using this model to page data in and out of memory dramatically improved system performance, and laid the foundations for a multitude of other optimisations, including efficient caching, network configuration and search engine ranking. There are two key aspects of this early work that demonstrate the value of a 'human-like' approach to software engineering: the first is that locality is a product of the way people write programs, rather than being due to any underlying constraints of the computing system; the second is that it was discovered empirically – paging algorithms were ineffective for many years before the principle of locality was discovered through systematic observation of the programming process [2].

Advances in hardware, such as parallel processing, have yet to achieve their full potential, as we struggle to translate serial human-written programs onto distributed architectures. Automated programming offer a means to reduce and repair errors, but even with machine-written programs, human input to a system means a bottleneck will always remain. As we move into the era of quantum computing and beyond, a true understanding of how our minds map themselves onto the machines we create is a vital component of achieving a step change in the creation and performance of software.

References

1 Caroline Jay and Robert Haines and Markel Vigo and Nicolas Matentzoglu and Robert Stevens and Jonathan Boyle and Alan Davies and Chiara Del Vescovo and Nicolas Gruel and Anja Le Blanc and David Mawdsley and Dale Mellor and Eleni Mikroyannidi and Richard Rollins and Andrew Rowley and Julio Vega Hernandez. *Identifying the challenges of code/theory translation: report from the Code/Theory 2017 workshop*. DOI: 10.3897/rio.3.e13236. Research Ideas and Outcomes, Pensoft Publishers, 2017.

2 Peter Denning. *The locality principle*. Communications of the ACM, 2007.

3.13 (Human-Like) Anecdotes with IBM's Watson/Bluemix Services

Kai-Uwe Kühnberger (Universität Osnabrück, DE)

Joint work of Kai-Uwe Kühnberger, Gordon Pipa, Students

I will give an overview of some activities in our institute in order to exemplify the idea of cognitively inspired AI. In particular, I will focus on cooperation projects with IBM, the company that coined the term "cognitive computing". Several terms are currently used for this or similar research endeavors, e.g. human-like AI, human-level AI, AI on a human scale, Artificial General Intelligence, Cognitive Computing, or Cognitive Services. They do not mean the same, but they are somehow related to each other. In my talk, I will briefly sketch applications like a smart city guide, smart farming, and flue prediction that were partially built on IBM's Bluemix services. Additionally I will sketch the idea to build an embodied e-tutor, a project that is currently ongoing. A focus will be to stress aspects that are easy (or straightforward) to implement and aspects that are rather difficult to realize.

3.14 From Turing to Deep Learning: Explaining AI through neurons and symbols

Luis C. Lamb (Federal University of Rio Grande do Sul, BR)

Learning and reasoning have been the subject of great research interest since the dawn of Artificial Intelligence [8]. In the 1950s, Turing already described principles for neural computation, machine learning and formal computational reasoning. Over the last decades, AI research has been questioned and praised on several occasions. However, recent developments in neural learning, in particular deep neural networks, have greatly impacted not only the academic research community, but also have been recognized as a key technology by both the computing industry and popular media channels [7]. In this talk, I presented an overview of the evolution of machine learning in AI, with particular attention to developments and efforts towards integrating (deep) learning and reasoning methods into an unified explainable foundation [2, 5]. I concluded showing that advances in AI, in particular neural-symbolic computation lead to the construction of rich computing systems that integrate neural learning, temporal and cognitive reasoning with applications in several areas [1, 3, 4, 5, 6].

References

1 Rafael V. Borges, Artur d'Avila Garcez, Luis C. Lamb. Learning and Representing Temporal Knowledge in Recurrent Networks. *IEEE Transactions on Neural Networks*, 22(12):2409–2421, Dec. 2011. IEEE Press.
2 A. S. d'Avila Garcez, Marco Gori, Pascal Hitzler, Luís C. Lamb: Neural-Symbolic Learning and Reasoning (Dagstuhl Seminar 14381). Dagstuhl Reports 4(9): 50–84, 2015.
3 A.S. d'Avila Garcez, Luis C. Lamb, and Dov Gabbay. Connectionist Modal Logics: Representing Modalities in Neural Networks. *Theoretical Computer Science*, 371(1–2):34–53, 2007
4 Artur S. d'Avila Garcez, Luis C. Lamb, Dov M. Gabbay. *Neural-Symbolic Cognitive Reasoning*. Springer, 2009.
5 A. d'Avila Garcez, Tarek R. Besold, Luc De Raedt, Peter Földiak, Pascal Hitzler, Thomas Icard, Kai-Uwe Kühnberger, Luis C. Lamb, Risto Miikkulainen, Daniel L. Silver. *Neural-Symbolic Learning and Reasoning: Contributions and Challenges*. Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches. Stanford Univ., March 2015, pp. 18–21, AAAI Press, 2015.
6 H.L. de Penning, Artur S. d'Avila Garcez, Luis C. Lamb, John-Jules Ch. Meyer: *A Neural-Symbolic Cognitive Agent for Online Learning and Reasoning*. Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence, IJCAI 2011. Pages 1653–1658, 2011.
7 Y. LeCun, Y. Bengio, G. Hinton. Deep Learning. *Nature*, 521 (7553): 436–444, 2015.
8 Alan Turing. Computing Machinery and Intelligence, *Mind*, LIX (236): 433–460, 1950.

3.15 Interacting Conceptual Spaces

Martha Lewis (University of Oxford, GB)

Joint work of Joe Bolt, Bob Coecke, Fabrizio Genovese, Martha Lewis, Dan Marsden, Robin Piedeleu
Main reference J. Bolt, B. Coecke, F. Genovese, M. Lewis, D. Marsden, R. Piedeleu, "Interacting Conceptual Spaces I : Grammatical Composition of Concepts", arXiv:1703.08314v2 [cs.LO], 2017.
URL https://arxiv.org/abs/1703.08314v2

How should we represent concepts and how can they be composed to form new concepts, phrases and sentences? The categorical compositional distributional programme of [2] successfully integrates two fundamental aspects of language meaning: firstly, the symbolic approach in which meanings of words compose to form larger units; and secondly, the distributional approach where word meanings are derived automatically from text corpora. These two approaches are unified by the key insight that each approach carries the same abstract structure, formalized using category theory.

The abstract framework of the categorical compositional scheme is actually broader in scope than natural language applications. It can be applied in other settings in which we wish to compose meanings in a principled manner, guided by structure. The outline of the general programme is as follows [1]:

1. a. Choose a compositional structure, such as a pregroup or combinatory categorial grammar.
   b. Interpret this structure as a category, the *grammar category*.
2. a. Choose or craft appropriate meaning or concept spaces, such as vector spaces, density matrices, or conceptual spaces.
   b. Organize these spaces into a category, the *semantics category*, with the same abstract structure as the grammar category.
3. Interpret the compositional structure of the grammar category in the semantics category via a functor preserving the necessary structure.
4. This functor then maps type reductions in the grammar category onto algorithms for composing meanings in the semantics category.

In this talk I describe how this programme can be applied to a number of different meaning spaces, including vector spaces, density matrices, conceptual spaces [3], and neural meaning spaces [4] and describe how these fit into the general programme described above. There are also choices for the syntactic side. I give examples of different grammars that can be used, and discuss how grammars might be built based on the semantic category chosen.

Acknowledgements: This work was partially funded by AFSOR grant "Algorithmic and Logical Aspects when Composing Meanings", the FQXi grant "Categorical Compositional Physics", and EPSRC.

References

1 Joe Bolt, Bob Coecke, Fabrizio Genovese, Martha Lewis, Dan Marsden, and Robin Piedeleu. Interacting conceptual spaces i: Grammatical composition of concepts. *arXiv preprint arXiv:1703.08314*, 2017.
2 B. Coecke, M. Sadrzadeh, and S. Clark. Mathematical foundations for a compositional distributional model of meaning. *Linguistic Analysis*, 36:345–384, 2010.
3 P. Gärdenfors. *Conceptual spaces: The geometry of thought.* The MIT Press, 2004.
4 Paul Smolensky. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, 46(1-2):159–216, 1990.

3.16 Deep Learning with Symbols

Daniel L. Silver (Acadia University – Wolfville, CA)

Joint work of Daniel L. Silver, Moh. Shameer Igball and Ahmed Galila at Acadia University, Wolfville, NS, Canada (2015-2017)
Main reference M. S. Iqbal, D. Silver, "A Scalable Unsupervised Deep Multimodal Learning System", in Proc. of the 29th Int'l FLAIRS Conf. (FLAIRS 2016), pp. 50–55, AAAI Press, 2016.
URL http://www.aaai.org/ocs/index.php/FLAIRS/FLAIRS16/paper/view/12928
Main reference Y. Bengio, "Evolving culture vs local minima', arXiv:1203.2990v2 [cs.LG], 2012.
URL https://arxiv.org/abs/1203.2990v2
Main reference I. Goodfellow, Y. Bengio, A. Courville, "Deep learning", in Adaptive computation and machine learning, MIT Press, 2016.
URL http://www.deeplearningbook.org/

We present work on the development of (1) a multimodal learning system that uses three sensor/motor channels and a symbolic channel to learn handwritten digits and (2) a system that learns to add handwritten digits aided by symbols of those digits. The conjecture is that symbols are predominately external communication tools that allow agents to share complex noisy concepts that help avoid local minimum in model development. They do so by providing secondary modeling tasks that provide beneficial inductive bias during learning, therefore reducing the number of examples required to accurate learn a new concept. The multimodal learning system uses a visual, auditory and motor channel constructed from stacks of RBMs. These stacks are brought together and connected to one large (2000 neurons) top level RBM and trained using contrastive divergence and a novel back-fitting algorithm. A symbolic channel was also added to display the concept that the activations in the top layer is representing. This addition was found to significantly improve the models performance on all channels. This supports the idea that classification of sensory input as a secondary task provides beneficial inductive bias. In the second network we test this idea by learning to add noisy MNIST digits with and without binary symbolic values at the input and output of a deep BP network. We propose that the models that are augmented with symbols, particularly at the output develop better models. Early experiments support this proposal.

3.17 Grounded Language Processing and Learning

Michael Spranger (Sony CSL – Tokyo, JP)

The talk discusses language learning from the viewpoint developmental robotics. We are interested in investigating how robots can learn to communicate starting from basic sensorimotor capacities through various stages of communicative proficiency: gesture [1], lexicon [3, 4] and grammar [2]. The focus of these experiments is to build grounded models of language processing but also how communication systems can be acquired. We discuss various algorithmic approaches from symbolic to probablistic to neural networks and highlight their current performance and research issues.

References

1 Spranger, Michael and Steels, Luc. *Discovering communication through ontogenetic ritualisation*. Proc. Joint IEEE International Conferences on Development and Learning and Epigenetic Robotics (ICDL-Epirob), pages 14–19, IEEE Press, 2014. DOI:10.1109/DEVLRN.2014.6982948.

2 Spranger, Michael and Steels, Luc. *Co-Acquisition of Syntax and Semantics – An Investigation in Spatial Language*. Proc. IJCAI'15: Proceedings of the 24th international joint conference on Artificial intelligence, pages 1909–1915, AAAI Press, 2015.
3 Spranger, Michael and Beuls, Katrien. *Referential Uncertainty and Word Learning in High-dimensional, Continuous Meaning Spaces*. Proc. Joint IEEE International Conferences on Development and Learning and Epigenetic Robotics (ICDL-Epirob), IEEE Press, 2016.
4 Nevens, Jens and Spranger, Michael. *Computational Models of Tutor Feedback in Language Acquisition*. Proc. Joint IEEE International Conferences on Development and Learning and Epigenetic Robotics (ICDL-Epirob), IEEE Press, 2017.

3.18 Are 'logically naive' people logically naive?

Keith Stenning (University of Edinburgh, GB)

Joint work of Theodora Achourioti, Keith Stenning
Main reference T. Achourioti, A. J. B. Fugard, K. Stenning, "The empirical study of norms is just what we are missing", Frontiers in Psychology, 5:1159, 2014.
URL https://doi.org/10.3389/fpsyg.2014.01159

The categorial syllogism is the oldest topic of the psychology of reasoning with the first paper by Störring [2]. Experimenters have agreed that their subjects are trying to do classical logical reasoning, and that they are, on the whole, doing it very badly. Recently, many have rejected the interest of the syllogism task and argued for probability as the appropriate normative framework for human reasoning. Stenning and Yule [1] showed that it is in fact hard to tell the difference in this task's data between a subject reasoning perfectly by classical logic, and one reasoning with tactically refined nonmonotonic methods; and produced evidence that most subjects can be construed as doing 'preferred model' reasoning in Logic Programming i.e. cooperative nonmonotonic discourse processing.

The multiple logics framework of Stenning & van Lambalgen shows that cooperative expository discourse is well modelled in logic programming, and has been construed as proposing that human reasoning is conducted in LP. Here I will present evidence that 'logically naive' subjects are perfectly capable of trying to reason classically if they are given a task that is clearly intended as the conduct of an adversarial dispute: namely providing counterexamples against invalid inferences proposed by an untrustworthy competitive agent. In short, their counterexamples establish that their understanding of classical conditionals in this situations is material implication, because their understanding of validity in this situation is of the absence of counterexamples. They are by no means tactically adept, but counter modelling in these conditions is a hard task, and induces a completely different attitude than the standard 'draw-a-conclusion' task.

So, at least two logics are needed to model naive human reasoning. This raises crucial questions for the relation between symbolic and sub-symbolic computation. Propositional LP is neurally implementable and highly tractable for searching large knowledgebases: an ideal logic for modelling implicit reasoning (System 1) because its search does not require supervision. Classical logic provided many of the insights into intractable computation, though the tiny syllogistic fragment is tractable. The real questions about the psychological implementation of classical logic are perhaps about the 'theorem prover' that has to guide reasoning, and its implementation.

So, at least two logics are needed to modelling naive human reasoning. This raises crucial questions for the relation between symbolic and sub-symbolic computation.

References

1 Stenning, K. and Yule, P. Image and language in human reasoning: a syllogistic illustration. *Cognitive Psychology*, 34:109–159, 1997.
2 Störring, G. Experimentelle untersuchungen über einfache Schlussprozesse, *Archiv für die Gesamte Psychologie*, 11:1–127, 1908.

3.19 Logic Programming as a framework for understanding human computation

Keith Stenning (University of Edinburgh, GB)

Main reference K. Stenning, M. van Lambalgen, "Human reasoning and Cognitive Science", MIT Press, 2008.
URL http://mitpress.mit.edu/books/human-reasoning-and-cognitive-science

Logic Programming (LP) has been used for analysing discourse and human reasoning with some success. The propositional form of LP is neurally implementable [1, 2]. This talk will explore the intersection this overlap between AI, Computer and Cognitive Science offers, raising some issues, from the cognitive side of the fence.

I will introduce four of the many issues that arise, chosen for possible relevance to neural-symbolic computation: 1) Computational tractability and questions of scaling-up; 2) Distinguishing reasoning and learning in psychology and computer science; 3) The relation of extensional to intensional reasoning; and 4) Qualitatively distinct kinds of uncertainty (especially relations between analyses in LP and probability). If discussion does not detain us, the talk will finish with an example of a current project at the nitty-gritty level: an analysis of human semantic memory in the service of nonmonotonic reasoning.

References

1 M. van Lambalgen, F. Hamm, *The Proper Treatment of Events*. Blackwell, Oxford and Boston, 2004.
2 Stenning, K. and van Lambalgen, M. *Human reasoning and Cognitive Science*. MIT Press, Cambridge, MA, 2008.

3.20 Structured Computer Organization of the Human Mind

Niels A. Taatgen (University of Groningen, NL)

Main reference N. A. Taatgen, "The nature and transfer of cognitive skills", Psychological Review, 120(3):439–471, 2013.
URL http://psycnet.apa.org/doi/10.1037/a0033138

Cognitive architectures aim to bridge the gap between the brain and behavior, providing a formal level of description that provides a basis for rigorous theories of behavior. Architectures typically operate on a single level of abstraction. In my talk I argued that we need multiple levels of abstraction, each with their own formalisms and learning mechanisms. Each of these levels should be to explain the abstraction level above, creating a reductionist hierarchy of theories that can bridge the gap, not with a single formalism, but with several.

3.21 Embodied neuro-symbolic computation

Serge Thill (University of Skövde, SE)

This talk covers two distinct reasons I am interested in neuro-symbolic integration. The first is because one is creating a cognitive system acting in the real world and carrying out some non-trivial task. In such systems, one often encounters problems that are better addressed with subsymbolic solutions such as neural networks and others for which symbolic approaches are more amenable. Here, I present the new European project Dreams4Cars (Grant agreement nr 731593) as one example of such a case.

The second reason is because it is a way to test representationalist embodied models of cognition: if embodied theories are right, then the sensorimotor experience underlying the formation of human concepts is critical for those concepts, and fundamentally shapes the computations that use them. An example effect one might expect in such a case is that the differences in sensorimotor experience between two agents might affect the degree to which they can communicate about concepts, or operations possible on concepts. It is therefore interesting to build a model of symbol grounding in a "rich" sensorimotor experience, and this cannot ignore the subsymbolic neural levels (since neural computations are fundamentally shaped by the constraints of biology, and possibly the specific needs of the organism) nor the details of the human sensorimotor experience (which, for example, includes interoceptive features).

The present talk attempts an overview of interesting issues in both directions.

3.22 Overview of neuro-symbolic processing in Neural Blackboard Architectures

Frank Van der Velde (University of Twente)

Joint work of Frank Van der Velde, Marc de Kamps
Main reference F. van der Velde, M. de Kamps, "Neural blackboard architectures of combinatorial structures in cognition", Behavioral and Brain Sciences, 29:37–70, 2006.
URL https://doi.org/10.1017/S0140525X06009022

Neural Blackboard Architectures (NBAs) aim to account for and simulate combinatorial structures and processing in a neuronal manner. Examples of these include sentential structures and reasoning with relations, but NBAs also include perceptual processing (vision). The focus in this presentation will be on sentential structures, both in relation to human processing and as a computational architecture for parallel processing. NBAs satisfy a number of constraints we assume to exist in neural and (human) cognitive processing. These include the assumption that representations (e.g. of "words") are grounded 'in situ' representations, also in combinatorial structures. NBAs provide "connection paths" between such representations, as a basis for generating behavior. Representations are always content-addressable in the underlying connection structures, also when they are a part of combinatorial structures such as sentences. Simulations with NBAs will be illustrated with incremental sentence processing and dynamical competitions, which can account for ambiguity resolution and garden path effects as found in human behavior.

References

1 van der Velde, F. and de Kamps, M. *From knowing what to knowing where: Modeling object-based attention with feedback disinhibition of activation*. Journal of Cognitive Neuroscience, 13(4), 479–491, 2001.
2 van der Velde, F. and de Kamps, M. *Learning of control in a neural architecture of grounded language processing*. Cognitive Systems Research, 11, 93–107, 2010.
3 van der Velde, F. and de Kamps, M. *Compositional connectionist structures based on in situ grounded representations*. Connection Science, 23, 97–107, 2011.
4 van der Velde, F. *Communication, concepts and grounding*. Neural Networks, 62, 112–117, 2015.
5 van der Velde, F. and de Kamps, M. *The necessity of connection structures in neural models of variable binding*. Cognitive Neurodynamics, 9, 359–370, 2015.
6 van der Velde, F. and de Kamps, M. *Combinatorial structures and processing in neural blackboard architectures*. In Proceedings of the Workshop on Cognitive Computation: Integrating Neural and Symbolic Approaches (CoCo@NIPS 2015), eds T. R. Besold, A. d'Avila Garcez, G. F. Marcus, and R. Miikkulainen (Montreal), (pp. 1–9). CEUR Workshop Proceedings, Vol. 1583, 2015.
7 van der Velde, F. *Concepts and relations in neurally inspired in situ concept-based computing*. Frontiers in Neurorobotics. 10:4, 2016.

3.23 Neural Computing with Signals and Symbols for Music Analysis

Tillman Weyde (City, University of London, GB)

Understanding music with computational methods is a challenging but rewarding task, as music is an essential and universal feature of humanity. Music information retrieval research develops methods that extract information from music data at different levels, increasingly using neural networks. The different representation levels of music as audio signal, notes, harmony, form etc, as well as emotional and cultural aspects, naturally create a need to connect unstructured and symbolic representations. In this talk, examples are discussed to demonstrate trends, achievements, and challenges in the development of human-like music computation.

References

1 Andreas Jansson and Eric Humphrey and Nicola Montecchio and Rachel Bittner and Aparna Kumar and Tillman Weyde. *Singing Voice Separation with Deep U-Net Convolutional Networks*. In Proc. of the International Society for Music Information Retrieval, ISMIR 2017, Suzhou, China, 2017 (in press).
2 Siddharth Sigtia and Emmanouil Benetos Nicolas Boulanger-Lewandowski and Tillman Weyde and Artur d'Avila Garcez and Simon Dixon. *A hybrid recurrent neural network for music transcription*. In Proc. of the International Conference on Acoustics, Speech and Signal Processing, ICASSP, Brisbane, Australia, 2015.

3 Reinier de Valk and Tillman Weyde. *Bringing 'Musicque into the tableture': machine-learning models for polyphonic transcription of 16th-century lute tablature*. Early Music 43.4 (2015): 563–576.

4 Tilman Weyde and Klaus Dalinghaus. *A neuro-fuzzy system for sequence alignment on two levels.* Mathware & Soft Computing 11.3 (2004):197–210.

4 Working groups

4.1 Logic Tensor Networks Hackathon

Lucas Bechberger (Universität Osnabrück, DE), Artur d'Avila Garcez (City, University of London, GB), Raquel Garrido Alhama (University of Amsterdam, NL), Marco Gori (University of Siena, IT), Luciano Serafini (Bruno Kessler Foundation – Trento, IT), Michael Spranger (Sony CSL – Tokyo, JP), and Tillman Weyde (City, University of London, GB)

Logic Tensor Networks (LTNs) [1] provide a neural network based solution for combining symbolic knowledge (encoded as first-order logic relations) and subsymbolic knowledge (represented as feature vectors). Each object is described as a vector of real numbers which together with relational symbolic knowledge get translated into soft and hard constraints on the subsymbolic level (implemented as a tensor network). The network then learns to approximate a solution to the constraint-optimization problem called best satisfiability when faced with new data. LTNs enable therefore relational knowledge infusion into deep networks, and knowledge completion and distilling through querying the networks.

In this hackathon, we investigated the application of the LTN framework to the domain of word embeddings. A word embedding such as word2vec [2] supplies for each word in language a so called "embedding", i.e. a vector in a high-dimensional space. These vectors are usually learned based on co-occurence patterns in a large text corpus. Somewhat surprisingly, the vectors of semantically related words tend to have a relatively low cosine distance. One can even go further and perform some arithmetic on the vectors (the vector resulting from "king – man + woman" is quite similar to the vector for "queen").

In our application scenario, we tried to generate new word vectors for previously unknown words based on some logical constraints that could for instance come from a dictionary entry or from an ontology. Our showcase looked as follows: we assumed that there is no vector embedding for the word "zebra", but there are embeddings for "horse" and for "black and white stripes". By defining an object as a zebra if and only if it is a horse and has black and white stripes, we wanted to derive a word embedding for "zebra". For a first proof-of-concept implementation, we used only a two-dimensional space with the following vectors for "horse" and "black and white": horse = (1.0, 0.0), black and white = (0.0, 1.0). We also added vectors for some additional unrelated words (e.g. cow). When running the LTN framework multiple times, we typically obtained vectors in the first quadrant, i.e. vectors for which both coordinate entires were greater than zero, as should be expected.

This was found to serve as a good example for introducing LTNs to academics with distinct backgrounds (e.g. symbolic logic/AI and statistical Machine Learning) and to help them get familiarized with the Tensorflow implementation of LTN, which is available at https://github.com/LucianoSerafini/dagstuhl_hackaton_LTN.

Following results reported in [1] of the application of LTNs to semantic image interpretation (c.f. https://uk.arxiv.org/abs/1705.08968) which can improve the performance of state-of-the-art image detection with convolutional networks, we are optimistic that LTNs can also help improve results in the natural language processing area by connecting word embeddings to symbolic knowledge from ontologies.

References

1 Ivan Donadello, Luciano Serafini, and Artur d'Avila Garcez. *Logic Tensor Networks for Semantic Image Interpretation*. Proceedings of the 26th International Joint Conference on Artificial Intelligence, 2017

2 Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeff Dean. *Distributed Representations of Words and Phrases and Their Compositionality*. Advances in Neural Information Processing Systems, 2013.

4.2 Explainability Hackathon

Tarek R. Besold (Universität Bremen, DE), Derek Doran (Wright State University – Dayton, US), Pascal Hitzler (Wright State University – Dayton, US), Caroline Jay (University of Manchester, GB), Oliver Kutz (Free University of Bozen-Bolzano, IT), and Sarah Schulz (Universität Stuttgart, DE)

Quite frequently demands for better comprehensible and explainable Artificial Intelligence (AI) and Machine Learning (ML) systems are being put forward. The explainability hackathon was intended to shed light on notions such as "comprehensibility" and "explanation" in the context of AI and ML, working towards a better understanding of what an explanation is when talking about intelligent systems, what it means to comprehend a system and its behavior, and how human-machine interaction can take these dimensions into account.

The hackathon compared notions of comprehension and explanation as used in different fields, ranging from philosophy to cognitive science, and mapped the findings back into AI and ML, aiming to establish generic types of systems. Three conceptual categories emerged, namely systems offering no insight into their input/output mechanisms; systems where users can mathematically analyze these mechanisms; and systems emitting symbols that allow a user to draw an intuitive explanation of how a conclusion is reached.

Following the Dagstuhl seminar, this initial line of work has been further expanded into a position paper titled "What Does Explainable AI Really Mean? A New Conceptualization of Perspectives" which has been submitted to the AI*IA 2017 workshop on Comprehensibility and Explanation in AI and ML (CEx). In that paper, additionally efforts towards systems of a fourth type are encouraged: explainable systems, where automated reasoning is central, yet missing from much of the current work.

4.3 Deep Learning with Symbols Hackathon

Daniel L. Silver (Acadia University – Wolfville, CA), James Christopher Davidson (Google Inc. – Mountain View, US), Dieuwke Hupkes (University of Amsterdam, NL), Isaac Noble (Playground Global Inc., US), and Katja Seeliger (Radboud University Nijmegen, NL)

Background: Recent work by Igbal and Silver (FLAIRS-2016 Best Paper Award) [1] inspired by work by Hinton et al. [2] and Srivastava et al. [3] has shown that it is possible to develop a multimodal deep learning system for learning a noisy handwritten digits using four sensor/motor channels (visual, audio, robotic, and symbolic) and an associative layer that ties all channels together. After training, the presentation of a digit (sound, image, drawing) at the visible nodes of the model activates all other channels to create their associated reconstruction at their respective visible nodes. Each channel provides additional information that helps the other channels more accurately reconstruct the output at their visible nodes. The symbolic channel outputs the cleanest and clearest signal as to what digit the multimodal deep learning system is "thinking" of given input on another channel. The symbolic channel also provides the cleanest and clearest input to assist other channels to generate the correct reconstructions at their visible nodes. This led us to a paper by Yoshua Bengio [4] that discusses the value of symbols (ie. language) in helping individuals to learn concepts (like "cat") better without having to see all possible examples of that concept. This has a profound impact upon the development of our culture and the human species.

Purpose: To create a deep recurrent neural network architecture that can add two or more noise MNIST digits in a row and produce the correct symbolic response at the output. The background provided above inspired us to consider a learning agent that learns to perform a mathematic operation using two noising channels but which can (at times) also receive concise information on a symbolic channel about the data on the noisy channels. We chose to investigate a system that could learn to add a sequence of handwritten digits aided (on occasion) by symbols of those digits at additional inputs to the same network. The conjecture is that symbols are predominately external communication tools that allow agents to share complex noisy concepts that help avoid local minimum in model development. They do so by providing secondary tasks that provide beneficial inductive bias during learning, therefore reducing the number of examples required to accurate learn a new concept. Our intention was to test this idea by developing a deep recurrent neural network using Keras/Tensorflow to add a sequence of noisy MNIST digits with and without binary symbolic values at the inputs and outputs of the network.

Progress: With much coding work by Dieuwke and Katja and support from Danny, James, and Isaac we were able to develop a Keras/Tensorflow-based LSTM recurrent network that could accept two sequential MNIST digits and output a symbolic value (two sets of 10 binary nodes) indicating the estimated result of the addition. Some of us intend to pursue this direction further using similar methods.

References

1 M.S. Iqbal and D. Silver. A Scalable Unsupervised Deep Multimodal Learning System. Proceedings of the Twenty-Ninth International FLAIRS Conference (FLAIRS-29), Key Largo, FL, AAAI Press, pp.50-55, 2016.
2 G.E. Hinton, S. Osindero and Y.-W. Teh. A fast learning algorithm for deep belief nets. Neural computation 18(7):1527– 1554, 2016.

3 N. Srivastava and R. R. Salakhutdinov. Multimodal learning with deep Boltzmann machines. In Advances in neural information processing systems, 2222–2230, 2012.
4 Yoshua Bengio. Evolving Culture vs Local Minima, March, 2012, https://arxiv.org/abs/1203.2990

4.4 Nengo/ACT-R/PRIMs/Cognitive Architectures Hackathon

Niels A. Taatgen (University of Groningen, NL), Jelmer Borst (University of Groningen, NL), Marc de Kamps (University of Leeds, GB), Martha Lewis (University of Oxford, GB), Serge Thill (University of Skövde, SE), and Frank Van der Velde (University of Twente)

Cognitive architectures aim to explain and simulate various or all aspects of human cognition and intelligence. Some of these architectures model cognition using symbolic representations, or representations close to symbolic, whereas others focus on neural simulations. A challenge for future progress is to integrate both approaches in order to benefit from accomplishments of both. In this Hackathon we discussed this integration between ACT-R and PRIMs (symbolic) on the one hand, and Nengo and the Neural Blackboard architecture on the other hand. In particular, we discussed how the smallest elements of skill proposed by the PRIMs architecture (primitive operations) can be implemented in a Basal Ganglia model in either Nengo or the Neural Blackboard architecture. As benchmark to test such models, we selected a set of monkey studies in which the monkey had to make controlled eye-movements. In addition, we talked about how conceptual spaces models can be viewed in the light of cognitive/neural architectures. To further develop this research, we agreed to organize a future workshop in Amsterdam in December 2017.

5 Open problems

5.1 Research Challenges

Luis C. Lamb (Federal University of Rio Grande do Sul, BR), Tarek R. Besold (Universität Bremen, DE), and Artur d'Avila Garcez (City, University of London, GB)

Recent developments in AI, in particular machine learning and deep networks demands explainability. In particular, humans need to trust AI systems and applications. Novel, state-of-the-art AI systems will have to provide explanations for their actions, behaviours and their impact on humanity. As mentioned above, this seminar consisted of both invited and contributed talks, joint group discussion sessions, and scientific hackathons. Several questions and research challenges were identified. One clear point made by the participants is that Human-Like Neural-Symbolic Computing is an effective methodology towards explainable AI [2]. As a result, Neural-Symbolic Computing clearly contributes towards answering questions that are both scientifically, ethically and methodologically relevant in the context of current Artificial Intelligence research.

In summary, the main research directions identified in the seminar are:

- The recent success of deep learning in computer vision and language processing [10], associated with the growing complexity of big data applications demands the development of *explainable AI models*. Neural-symbolic computation offers several methods and principles that contribute towards this aim. For instance, rule extraction mechanisms [3] and the integrated methodology for the development of neural-symbolic cognitive systems offers principled alternatives [1, 5, 6, 8, 9]. The concept of modularity in deep networks can be explained by using neural-symbolic methods. Modularity is suitable to knowledge extraction (as identified in [4, 5]) which can result in the development of effective knowledge extraction methods towards explainable AI.
- The development of Cognitive Architectures capable of simulating and explaining aspects of human cognition also remains an important research endeavour [11]. Some cognitive architectures typically consider symbolic representations, whereas others employ neural simulations. The integration of these models remains a challenge and there are benefits on integrating the accomplishments of both paradigms, as identified in the cognitive architectures Hackathon.
- Logic Tensor Networks (LTNs) [7] provide a model that integrates symbolic knowledge (encoded as first-order logic relations) and subsymbolic knowledge (represented as feature vectors). LTNs enable the representation of relational knowledge infusion into deep networks, and knowledge completion and distilling through querying the networks. There remains a number of challenges in integrating, explaining and computing symbolic knowledge in deep networks. Both LTNs [7] and Connectionist Modal and Temporal Logics [5, 6, 9] offer effective alternatives towards these research challenges.

References

1 Tarek R. Besold, Artur S. d'Avila Garcez, Keith Stenning, Leendert W. N. van der Torre, Michiel van Lambalgen: Reasoning in Non-probabilistic Uncertainty: Logic Programming and Neural-Symbolic Computing as Examples. *Minds and Machines*, 27(1): 37–77, 2017.

2 A. d'Avila Garcez, Tarek R. Besold, Luc De Raedt, Peter Földiak, Pascal Hitzler, Thomas Icard, Kai-Uwe Kühnberger, Luis C. Lamb, Risto Miikkulainen, Daniel L. Silver. Neural-Symbolic Learning and Reasoning: Contributions and Challenges. Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches. Stanford Univ., March 2015, pp. 18–21, AAAI press, 2015.

3 Artur S. d'Avila Garcez, Dov M. Gabbay, Krysia Broda. *Neural-Symbolic Learning System: Foundations and Applications*. Springer-Verlag, New York, Inc., USA, 2002.

4 A. S. d'Avila Garcez, Marco Gori, Pascal Hitzler, Luís C. Lamb: Neural-Symbolic Learning and Reasoning (Dagstuhl Seminar 14381). Dagstuhl Reports 4(9): 50–84, 2015.

5 Artur S. d'Avila Garcez and Luis C. Lamb. A Connectionist Computational Model for Epistemic and Temporal Reasoning. *Neural Computation*, 18(7):1711–1738, 2006.

6 Artur S. d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay. *Neural-Symbolic Cognitive Reasoning*. Springer Publishing Company, 2009. ISBN: 9783540732457.

7 Ivan Donadello, Luciano Serafini, and Artur d'Avila Garcez. *Logic Tensor Networks for Semantic Image Interpretation*. Proceedings of the 26th International Joint Conference on Artificial Intelligence, 2017.

8 Barbara Hammer, Pascal Hitzler (Eds.): Perspectives of Neural-Symbolic Integration. Studies in Computational Intelligence 77, Springer 2007.

9 Luis C. Lamb, R.V. Borges and A.S. d'Avila Garcez. *A connectionist cognitive model for temporal synchronisation and learning*. Proc. AAAI Conference on Artificial Intelligence AAAI-07, pages 827–832, AAAI Press, 2007.

10 Y. LeCun, Y. Bengio, G. Hinton. Deep Learning. *Nature*, 521 (7553): 436–444, 2015.

11 K. Stenning and M. van Lambalgen. *Human reasoning and Cognitive Science*. MIT Press, Cambridge, MA, 2008.

Participants

- Lucas Bechberger
Universität Osnabrück, DE
- Tarek R. Besold
Universität Bremen, DE
- Jelmer Borst
University of Groningen, NL
- Artur d'Avila Garcez
City, University of London, GB
- James Christopher Davidson
Google Inc. –
Mountain View, US
- Marc de Kamps
University of Leeds, GB
- Derek Doran
Wright State University –
Dayton, US
- Ulrich Furbach
Universität Koblenz-Landau, DE
- Raquel Garrido Alhama
University of Amsterdam, NL
- Marco Gori
University of Siena, IT
- Pascal Hitzler
Wright State University –
Dayton, US
- Dieuwke Hupkes
University of Amsterdam, NL
- Caroline Jay
University of Manchester, GB
- Kristian Kersting
TU Darmstadt, DE
- Kai-Uwe Kühnberger
Universität Osnabrück, DE
- Oliver Kutz
Free University of
Bozen-Bolzano, IT
- Luis C. Lamb
Federal University of
Rio Grande do Sul, BR
- Martha Lewis
University of Oxford, GB
- Isaac Noble
Playground Global Inc., US
- Sarah Schulz
Universität Stuttgart, DE
- Katja Seeliger
Radboud University
Nijmegen, NL
- Luciano Serafini
Bruno Kessler Foundation –
Trento, IT
- Daniel L. Silver
Acadia University –
Wolfville, CA
- Michael Spranger
Sony CSL – Tokyo, JP
- Keith Stenning
University of Edinburgh, GB
- Niels A. Taatgen
University of Groningen, NL
- Serge Thill
University of Skövde, SE
- Frank Van der Velde
University of Twente, NL
- Tillman Weyde
City, University of London, GB

Report from Dagstuhl Seminar 17201

Formal Synthesis of Cyber-Physical Systems

Edited by

Calin A. Belta[1], Rupak Majumdar[2], Majid Zamani[3], and Matthias Rungger[4]

1 **Boston University – Brookline, US,** cbelta@bu.edu
2 **MPI-SWS – Kaiserslautern, DE,** rupak@mpi-sws.org
3 **TU München, DE,** zamani@tum.de
4 **TU München, DE,** matthias.rungger@tum.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17201 "Formal Synthesis of Cyber-Physical Systems." Formal synthesis is the application of algorithmic techniques based on automata and logic to the design of controllers for hybrid systems in which continuous components interact with discrete ones. The Dagstuhl seminar brought together researchers from control theory and from computer science to discuss the state-of-the-art and current challenges in the field.

Seminar May 14–19, 2017 – http://www.dagstuhl.de/17201

1998 ACM Subject Classification I.2.8 Problem Solving, Control Methods, and Search—Control theory; I.2.2 Automatic Programming; F.3.1 Specifying and Verifying and Reasoning about Programs

Keywords and phrases Cyber-physical systems, formal synthesis, reactive synthesis, discrete event systems, dynamical systems and control

Digital Object Identifier 10.4230/DagRep.7.5.84

1 Executive Summary

Calin A. Belta
Rupak Majumdar
Majid Zamani
Matthias Rungger

Cyber-Physical Systems (CPS) are complex systems resulting from intricate interaction of discrete computational devices with the continuous physical plants. Within CPS, embedded control software plays a significant role by monitoring and adjusting several physical variables, e.g. temperature, velocity, pressure, density, and so on, through feedback loops where physical processes interact with computational devices. Recent advances in computation, storage, and networking have made tremendous advances in hardware and system platforms for CPS. With this growing trend in computational devices, embedded control software is becoming more and more ubiquitous in many safety-critical applications including automotive, aerospace, transportation systems, critical infrastructure, energy, robotics, healthcare, and many other domains. Unfortunately, the design of embedded control software nowadays is still based on ad-hoc solutions resulting in brittle and error-prone software, and very high verification and validation costs. In order to detect and eliminate design flaws and inevitable software bugs, a large portion of the design budget is consumed with validation and verification efforts, which are often lengthy. On the other hand, by changing the emphasis from verification to synthesis, it is possible to synthesize correct-by-design embedded control software for CPS while providing formal guarantees of correctness and preventing the need for costly post facto verification.

In recent years, there has been a lot of progress in designing automatic and correct-by-construction techniques for controller synthesis for interacting discrete and continuous systems. These new techniques have combined techniques from continuous control theory as well as from computer science. The focus of this seminar was to provide a state-of-the-art of this nascent but important field, and to describe challenges and opportunities for synthesis techniques to transition to the real world. By the nature of the topic, the participation at the seminar was inter-disciplinary, and consisted of computer scientists and control theorists, both from academia and from industry. Instead of a sequence of presentations of individual research results, the seminar was organized as a sequence of open discussions on topics of common interest to the participants, such as techniques for scalable controller synthesis, identification of application domains and recent success stories, compositionality and system design, end-to-end arguments about systems, as well as education and outreach.

This seminar benefitted the control as well as computer science communities by bridging the gap between many complementary concepts studied in each community. A more detailed survey on the topics of the seminar is in preparation.

Outcomes of the seminar

The seminar focused on the challenges in the application of formal synthesis techniques for automatic, correct-by-construction synthesis of CPS. The seminar had a total of 45 participants with a mix of computer scientists and control theorists.

Sessions

The seminar was organized as a sequence of open discussions led by one or two moderators. Each session had a scribe to note down the discussion. The scribe notes were shared with all participants, who added their comments or filled in more information. The updated notes were used to prepare the session summaries (in the next Section).

The following sessions were organized:

1. Application domains, success stories, and obstacles to adoption
2. Fundamental algorithmic and scalability challenges in formal synthesis
3. Tools and infrastructure
4. Education and outreach
5. Data-driven and search-driven approaches to synthesis
6. Compositionality in synthesis
7. Optimality and completeness
8. Synthesis of distributed protocols from scenarios and requirements.
9. Specification languages
10. Robustness and resiliency
11. Explainability and user interaction
12. End-to-end correctness
13. Formal synthesis challenges in robotics
14. Cyber-Security of CPS

2 Table of Contents

Executive Summary
Calin A. Belta, Rupak Majumdar, Majid Zamani, and Matthias Rungger 84

Summary of Discussions

Application Domains I: Robotics and Automotive Control
Calin Belta and Jyotirmoy Deshmukh 87

Application Domains II: Traffic Networks and Power Networks
Murat Arcak and Agung Julius 88

Fundamental algorithmic and scalability challenges in formal synthesis
Paulo Tabuada and Rupak Majumdar 89

Tools and infrastructure
Ruediger Ehlers and Matthias Rungger 89

Teaching and Outreach
Calin Belta and Stephane Lafortune 90

Data-driven synthesis and/or search-driven synthesis
Georgios Fainekos and Agung Julius 90

Compositionality in synthesis
Majid Zamani and Murat Arcak and Stavros Tripakis 91

Optimality and Correctness
Gunther Reissig and Antoine Girard 91

Robustness and resilience
Hadas Kress-Gazit and Pavithra Prabhakar 92

Explainability
Hadas Kress-Gazit 92

End-to-end correctness
Samarjit Chakraborty and Sanjoy Mitter 93

Application Domains III: Formal synthesis challenges in robotics
Hadas Kress-Gazit 94

Formal synthesis challenges in Cybersecurity
Paulo Tabuada and Ashutosh Trivedi 94

Participants 96

3 Summary of Discussions

3.1 Application Domains I: Robotics and Automotive Control

Calin Belta and Jyotirmoy Deshmukh

The session was scribed by Paulo Tabuada.

Applications are important as they identify the theoretical, scalability, and computational challenges. Thus, we started the seminar with an identification of domains where formal synthesis could be profitably applied, or has had some preliminary success. The following application domains were discussed:

- Automotive
- Robotics
- Power Networks, Power Systems, Smart Grids
- Transportation (Smart Cities, Traffic Networks)
- Avionics
- Medical Devices
- Biology
- Manufacturing
- Smart buildings

Based on the expertise present at the seminar, we focused on the automotive, robotics, traffic network, and power systems. As a general feature, it was pointed out that an advantage of formal methods is that it is easy to change specifications and get new controllers, whereas a manual method may need the designer to start from scratch.

In the robotics domain, we discussed how formal approaches can be reconciled with standard practice of "build, test, iterate." We discussed synthesis problems specific the robotics infrastructure such as the use of higher-level programming abstractions and declarative specifications, and problems arising out of software upgrades. We discussed the abstraction level at which synthesis should be carried out. For "simple" dynamics, it is common to abstract the continuous dynamics and present the synthesis problem as planning in a grid world. However, for more complicated tasks such as grasping or bipedal walking, control for the continuous dynamics is important. A related question is how to encode background knowledge into the synthesis problem definition? For example, in many robotics applications we may need a large set of predicates, but LTL synthesis algorithms work better when we have a small number of predicates. Existing logics may not model reactive systems.

In the *automotive control domain*, we discussed where formal synthesis would fit in the existing industrial design process. Model-based Controller Design using Simulink is a success story since as this entails a refinement-based design methodology. Skepticism towards formal control design in the industry could be attributed to lack of knowledge on using formal methods properly. For instance, it is often not clear at what level of abstraction should formal synthesis be applied. Formal methods can already provide solutions for some of the small problems that are solved during the design process. Formal methods help in writing contracts between components and guide the solution. Modest sub-problems can be handled with formal methods.

In the automotive context there are several hardware platforms and protocols for communication between distributed electronic control units. These protocols are well-defined and there are many tools for timing analysis at the hardware level. However, they are decoupled from Simulink and other tools that can analyze the desired global behavior; e.g., existing timing analysis tools do not care about stability and we may be solving a problem harder than needed. There seems to be a gap, and there may be a need for tools that act as glue between existing tools at various levels of abstraction. To make problems harder, these tools are used by different teams, potentially in different organizations.

The NSF expeditions program on Computer-Aided Program Engineering (ExCAPE) indicates that solving the verification problem can help you solve the synthesis problem using counter-example guided synthesis. An optimistic view is that a similar approach could be used for CPS. A challenge for CPS is that specifications are not articulated well, so solving the specification problem is a prerequisite for solving the verification problem!

It would be beneficial to create a road-map for the use of formal methods in industry. Benchmarks are crucial to understand application scenarios for formal synthesis. Synthesis competitions for CPS may be a good exercise, following existing competitions that cover reachability analysis and falsification.

3.2 Application Domains II: Traffic Networks and Power Networks

Murat Arcak and Agung Julius

This session was scribed by Paulo Tabuada.

Recently, formal synthesis techniques have been applied to congestion control problems in traffic networks. Discrete-time piecewise affine (PWA) systems are widely used as models for traffic networks. There exist works on verification and synthesis for such systems. By exploiting the sparsity and the structure inherent in PWA models of traffic networks, abstractions can be constructed for fairly large systems (almost 60 dimensions). Congestion can be characterized as polyhedral sets and included as predicates in temporal logic formulas used as specifications.

Some concrete challenges in this domain were identified:

1. The connection between microscopic models (VSim) and the PWA density based models is not understood. There are other models at different layers. not clear what the formal connection is between these models.
2. Existing models do not explicitly capture V2V, V2I.
3. There is a need for probabilistic approaches to traffic networks. This would not improve scalability, but it will be more realistic.
4. Not clear how to integrate existing approaches focused on traffic light and meter control policies with autonomous vehicles.
5. Not clear how to integrate other control inputs, e.g., speed limits?
6. Existing approaches are not real time. If a road in closed or there is an accident, it is not clear how to adapt the controller

As with the previous application domains, there are several opportunities for formal synthesis in power networks, e.g., avoiding cascading failures in a power grid (Susuki et al, IEICE Trans. Fundamentals 2009). We discussed an approach based on nonlinear dynamics and circuit breakers. Changing the operating conditions of some generators might prevent the cascading failure. Metric temporal logic (MTL) was used as specification to prevent current peaks. This problems is a transient regime problem. It was identified that abstractions in this domain should capture the transient behavior.

These networks have many time-scales dynamics ranging from secs to minutes to hours-intervals. For formal verification there are different sources of adversaries. They are also very large. One way to address scalability would be to exploit possible symmetries in the network.

3.3 Fundamental algorithmic and scalability challenges in formal synthesis

Paulo Tabuada and Rupak Majumdar

This session was scribed by Calin Belta.

The synthesis problem has a long history both in computer science as well as in control theory. At an abstract level, the synthesis problem can be stated as follows: given a system and a specification, synthesize a controller that restricts the behavior of the system so as to satisfy the specification. When the system is given as a finite-state automaton and the specification given as a linear-time temporal logic formula, the problem can be solved algorithmically with time complexity that is doubly exponential in the length of the formula and polynomial in the size of the system. These results can also be used when the system is described by a differential equation. In this case, we first construct a finite-state abstraction of the differential equation and then use the results for finite-state automata. Unfortunately, existing abstraction techniques result in finite-state automata with a size that is exponential in the number of variables appearing in the differential equation. Given the high complexity of existing solutions for the synthesis problem, we discussed in this session several promising research directions to manage the high computational complexity, including incremental synthesis algorithms, how to leverage interactions with the user, robust strategies, and the use of machine learning techniques.

3.4 Tools and infrastructure

Ruediger Ehlers and Matthias Rungger

This session was scribed by Anne-Kathrin Schmuck

Different dimensions of existing tools for synthesizing controllers were discussed, ranging from the type of models and specifications they support, the type of guarantees they provide to formal correctness of the implementation. There was a general consensus that a set of exemplary instances of synthesis problems serving as benchmark problems are missing in the context of CPS. Additionally, benchmark examples should be advertised and made publicly available, e.g. at www.cps-vo.org. In order to foster the adaptation of formal methods tools in industry, using Simulink would be suitable candidate, as there are already automated techniques for translating Simulink to formal models. Finally, it was observed that software tools play an important role in education.

3.5 Teaching and Outreach

Calin Belta and Stephane Lafortune

This session was scribed by Manuel Mazo.

The participants discussed who to incorporate CPS concepts and synthesis techniques in undergraduate and graduate curricula in engineering and computer science. Existing courses and recently-published textbooks were identified. It was suggested to maintain a list of such courses and their syllabi on the CPS-VO. Industry needs should be considered in the development of such courses. An important resource on CPS education is the recent US NAE report titled "A 21st Century Cyber-Physical Systems Education". The participants also discussed outreach to industry. The availability of software tools that can be used by engineers in industry was identified as key to allow transfer of the technologies developed in the CPS synthesis community.

3.6 Data-driven synthesis and/or search-driven synthesis

Georgios Fainekos and Agung Julius

Model based search methods have proven their value in several practical control applications as well as in verification and falsification of CPS. For example, search based methods can be used to compute paths for autonomous vehicles, or to compute insulin infusion schedules, and to detect errors in Model Predictive Controllers (MPC) used in closed-loop artificial pancreas systems. System verification methods can also benefit by the utilization of machine learning techniques. In addition, there is a great opportunity to combine synthesis methods with machine learning techniques in order to control systems with unknown or uncertain models. Success has already been demonstrated in robotics applications where robots learn automata modeling the behavior of the environment and, then, using this information to plan and complete their missions. Along these lines, it may be possible to use data to learn system requirements, or to classify data, and then design controllers which satisfy the learned requirements. Another, approach would be to directly use already classified desirable and/or undesirable system behaviors along formal requirements for control synthesis.

A general future direction that came out of this workshop is to investigate what are the links between synthesis and learning, and how to combine the two. There are at least two ways one can think of combining the two:

1. Learning techniques for existing verification/synthesis problems: Here the emphasis is on using, adapting, or getting inspiration from, techniques from the broad machine learning field, in order to solve existing ("traditional") synthesis (and perhaps also verification?) problems.
2. New verification/synthesis problems, that adopt concepts from machine learning: Here, the idea is to come up with new and interesting problems (and then of course with solution methods for these problems). One example of a perhaps new such problem is the SSE problem (synthesis from spec and examples). Another example could be new variants of the Learning problem mentioned in the same section. New variants could be created by revisiting the notion of what it means for a learned system to "generalize well."

Two major challenges/open problems were identified during the session discussion:

- What are the right problems to solve? Learning in verification vs verification of learning.
- Where is the boundary of solvable and unsolvable computational problems when simultaneous learning and synthesis are concerned? For example, can we simultaneously learn the controller and the objective function from observations? How do we combine synthesis methods from high level specifications with reinforcement learning methods?

3.7 Compositionality in synthesis

Majid Zamani and Murat Arcak and Stavros Tripakis

This session was scribed by Matthias Rungger.

The main aim of this session is to investigate potential directions on compositional synthesis of controllers enforcing some global high-level specifications over interconnected cyber-physical systems.

In general, distributed reactive synthesis and supervisory control are undecidable but several approaches solve the problem for special cases. We discussed promising approaches to compositional synthesis of controllers, as well as their current limitations:

- There is some initial work on specification decomposition but it does not take into system composition structure. Specifically, can one exploit the structure and the sparsity of interconnection topology?
- There are promising results on discrete event systems on control design by abstracting components, composing abstractions, and abstracting them again, and then solving the resulting synthesis problem monolithically. It was suggested that the benchmarks from DES should be used for reactive synthesis.
- Synthesis from component libraries has been proposed for hardware and software; control theory needs to investigate this direction more.
- An important open direction is the unification of compositional rules in control theory (e.g. small-gain type reasoning) and computer science (e.g. assume-guarantee reasoning).
- An input-output interconnection may not be appropriate in those scenarios in which the dynamics of each component change due to interconnection; What is a good contract formalism for such systems?

3.8 Optimality and Correctness

Gunther Reissig and Antoine Girard

The first part of the session was dedicated to the synthesis of controllers optimizing quantitative objectives while enforcing some qualitative properties. Such optimal control problems appear naturally when a cost (e.g. energy consumption) is naturally associated with trajectories. Sometimes, it is possible to encode the qualitative requirements inside the quantitative objectives (as in linear quadratic regulation where stability is a consequence of finiteness of the quantitative objective). Synthesis of robust controllers can also be tackled through optimal control, e.g., by maximizing bounds on admissible disturbances or by optimizing the quantitative value of a temporal logic formula given by its quantitative semantics.

There exists a variety of possible cost functions for optimal control problems: discounted costs are often used in control theory, while mean payoff objectives (i.e., long term average cost) have been considered in computer science. These costs functions are actually complementary. Mean payoff objectives only optimize the asymptotic behavior of a system, while discounted costs focus mainly on the transient dynamics. While optimal control with discounted costs can be solved using classical dynamic programming, there are no known polynomial algorithms to solve mean payoff games. Actually, optimal strategies in mean payoff games may require infinite memory. Other types of costs functions may include positive and negative costs as in consumption or energy games.

The second part of the session focused on the problem of completeness. Asymptotic completeness in abstraction-based control holds when it is possible to approximate the true solution (e.g., maximal controller, value function) arbitrarily closely by refining sufficiently the abstraction. Some completeness results can be given for systems admitting bisimilar abstractions (e.g., controllable linear systems). For other systems, asymptotic completeness does not hold for all problems and is related to robustness and continuity of the value function with respect to the problem data. These conditions are hard to check a priori so a posteriori estimates of performance gaps and convergence rates may be useful. A promising direction is the development of synthesis approach that are anytime correct / asymptotically complete, allowing to obtain quickly a correct solution, which can be refined iteratively if needed and, which approaches the maximal / optimal solution asymptotically.

3.9 Robustness and resilience

Hadas Kress-Gazit and Pavithra Prabhakar

Robust synthesis corresponds to the synthesis of a controller that along with the model satisfies a specification, even in the presence of perturbations in the model, controller and specification. Such perturbations capture uncertainty/temporary environment violations, measurement/actuation errors, and robust semantics of properties described in temporal lgoics such as LTL/STL. The discussion consisted of different versions of robust synthesis problem that captured different relations between perturbations in the model, controller, and specification. In addition, metrics were identified as important in formulating the robust synthesis problem, and different metrics were discussed.

3.10 Explainability

Hadas Kress-Gazit

The usual outcome of formal synthesis is a synthesized controller. However, there are other artifacts that can be automatically generated that will make impact on the user. In this session, we discussed the different artifacts that can be generated to debug specifications, reason about controllers and facilitate the composition of the controller in a system, focusing on:

- Can formal synthesis be leveraged for explainability?
- What artifacts can be generated through the synthesis process to aid the user in understanding/debugging/validating the controller and the specification?

What can a tool return when the specification is unrealizable? One can return counterstrategies to present the reasons for unrealizability. This can happen through understandable presentation of the counterstrategy or suggestions of assumptions that would make the specification realizable or example executions for the environment that make the system fail. Tools can also return minimal specification revisions to make it realizable. Dually, tools can identify vacuous parts of specifications. There is work in the formal methods community on detecting vacuous specifications (in hardware verification) that can be leveraged to identify vacuity in specifications used to synthesize controllers. Additional possibilities are sensitivity analysis of counterexamples for identifying fragile areas in the control, grouping together counterexamples, monitors for runtime assumption violation from the specification,

We discussed how to produce understandable witnesses that “prove” the black-box controller obtained as the output of the synthesis tool. Witnesses could be “incimplete but understandable”: for example, through decision trees, skeletons, simulations, or other “understandable” representations of the controller.

3.11 End-to-end correctness

Samarjit Chakraborty and Sanjoy Mitter

The goal of this session is to discuss how to integrate different layers of the design stack in embedded control systems in order to facilitate end-to-end validation and synthesis. Current workflows are based on independent design and analysis of these layers, often involving assumptions at the higher modeling and algorithms development layers that do not hold true at the lower implementation layers. This results in significant ex post facto integration, testing and debugging efforts and pose major roadblocks in the validation, certification and automated synthesis of such systems. Partitioning the entire design problem into independent layers was necessary for conquering design complexity and has led to significant progress in general-purpose computing, where no assumptions on the applications were anyway possible. Hence, the focus —at all fronts such as compilers, operating systems and computer architecture— has always been on optimizing a general set of performance metrics such as (average case) execution time, memory footprint, communication speed and power consumption. However, when it comes to control algorithms, first, these secondary metrics do not serve the purpose adequately since they do not necessarily optimize the primary metrics such as stability, settling time or peak overshoot. Second, there is a rich variety of design methods and proof techniques available within control theory which are now restricted to the modeling and algorithms design layer and do not permeate into the implementation layers below. This brings up the following broad research questions/directions:

- How to develop a cross-layer framework for end-to-end design and validation of embedded control systems?

- How to incorporate implementation level details directly in the control algorithms design phase so that different implementation-level resource tradeoffs can be accounted for during controller design?
- Control/architecture co-synthesis? As a starting point, the "architecture" could be a scheduler or a protocol.
- Synthesis on the fly: Control algorithm (or at least its parameters), along with its implementation level parameters to change on the fly, in response to either (a) changes in the plant, or (b) changes in resource availability.
- Proof transformation techniques: Example, develop compilation techniques that will transform the proofs along with the models, so as to automatically generate a certificate of correctness along with the implementation.
- Develop tool chains.

3.12 Application Domains III: Formal synthesis challenges in robotics

Hadas Kress-Gazit

Since the discussion on application domains had not converged, there was a separate session on formal synthesis challenges in robotics. The main points discussed in this seesion were hierarchical approaches to formal synthesis, different abstraction formalisms (including open questions, such as an appropriate formalism for reconfigurable robots), and synthesis approaches to sensor/actuaror failures on-the-fly. It was suggested that synthesis can have an impact in robotics in providing explanable, robust, and formal plans which can be automatically integrated into the software infrastructure. We discussed the distinction between planning and synthesis, and the additional benefits synthesis can bring, e.g., by explicitly modeling reactivity and environment assumptions, by providing runtime monitors about environment assumptions, and guarantees w.r.t. a model.

3.13 Formal synthesis challenges in Cybersecurity

Paulo Tabuada and Ashutosh Trivedi

Security and privacy concerns are increasingly relevant in the CPS context, where a security attack can very easily have catastrophic consequences. Security attacks on control systems are no longer a theoretical possibility: there have been much publicized incidents in the recent past (e.g., Stuxnet). We discussed synthesis approaches and research challenges for cyber-security, including:

1. Synthesizing for side-channel (timing or space) vulnerabilities.
2. How do we measure tradeoff between security and the effort for mitigating the leakage-risk? In CPS settings, this can often be modeled as making the system unobservable. In Discrete Event Systems, security enforcement can be achieved through the notion of system (opacity), which involves hiding info about current state or initial state.

3. Security Enforcement method in control theory by checking if the system is invertible or by designing a controller to destroy surjectivity of the map.
4. Unique challenges for CPS security include the unacceptable overhead of standard cryptographic primitives, novel attack-surfaces (attacks on analog sensors, side channels in physical processes, algorithms interacting with physical environment, resource-usage info not in code)

Participants

- Erika Abraham
RWTH Aachen, DE
- Murat Arcak
University of California – Berkeley, US
- Calin A. Belta
Boston University – Brookline, US
- Ivana Cerná
Masaryk University – Brno, CZ
- Samarjit Chakraborty
TU München, DE
- Chih-Hong Cheng
fortiss GmbH – München, DE
- Jyotirmoy Deshmukh
Toyota Motors North America, US
- Rüdiger Ehlers
Universität Bremen, DE
- Martin Fabian
Chalmers UT – Göteborg, SE
- Georgios Fainekos
Arizona State University – Tempe, US
- Antoine Girard
CNRS – Gif sur Yvette, FR
- Anak Agung Julius
Rensselaer Polytechnic Institute, US
- Roozbeh Kianfar
Zenuity, SE
- Hadas Kress-Gazit
Cornell University – Ithaca, US
- Jan Kretinsky
TU München, DE
- Stéphane Lafortune
University of Michigan – Ann Arbor, US
- Jun Liu
University of Waterloo, CA
- Rupak Majumdar
MPI-SWS – Kaiserslautern, DE
- Manuel Mazo
TU Delft, NL
- Sanjoy K. Mitter
MIT – Cambridge, US
- Sahar Mohajerani
Chalmers UT – Göteborg, SE
- Necmiye Ozay
University of Michigan – Ann Arbor, US
- Ivan Papusha
Univ. of Texas at Austin, US
- Nir Piterman
University of Leicester, GB
- Pavithra Prabhakar
Kansas State University – Manhattan, US
- Gunther Reißig
Universität der Bundeswehr – München, DE
- Harald Ruess
fortiss GmbH – München, DE
- Matthias Rungger
TU München, DE
- Indranil Saha
Indian Institute of Technology Kanpur, IN
- Anne-Kathrin Schmuck
MPI-SWS – Kaiserslautern, DE
- Sadegh Soudjani
MPI-SWS – Kaiserslautern, DE
- Paulo Tabuada
University of California at Los Angeles, US
- Wolfgang Thomas
RWTH Aachen, DE
- Hazem Torfah
Universität des Saarlandes, DE
- Stavros Tripakis
University of California – Berkeley, US
- Ashutosh Trivedi
University of Colorado – Boulder, US
- Jana Tumova
KTH Royal Institute of Technology – Stockholm, SE
- Marcell Vazquez-Chanlatte
University of California – Berkeley, US
- Majid Zamani
TU München, DE

Report from Dagstuhl Seminar 17202

Challenges and Opportunities of User-Level File Systems for HPC

Edited by

Andrë Brinkmann[1], Kathryn Mohror[2], and Weikuan Yu[3]

1 **Johannes Gutenberg University Mainz, DE,** brinkman@uni-mainz.de
2 **LLNL – Livermore, US,** mohror1@llnl.gov
3 **Florida State University – Tallahassee, US,** yuw@cs.fsu.edu

Abstract

The performance gap between magnetic disks and data processing on HPC systems has become that huge that an efficient data processing can only be achieved by introducing non-volatile memory (NVRAM) as a new storage tier. Although the benefits of hierarchical storage have been adequately demonstrated to the point that the newest leadership class HPC systems will employ burst buffers, critical questions remain for supporting hierarchical storage systems, including: How should we present hierarchical storage systems to user applications, such that they are easy to use and that application code is portable across systems? How should we manage data movement through a storage hierarchy for best performance and resilience of data? How do the particular I/O use cases mandate the way we manage data? There have been many efforts to explore this space in the form of file systems, with increasingly more implemented at the user level. This is because it is relatively easy to swap in new, specialized user-level file systems for use by applications on a case-by-case basis, as opposed to the current mainstream approach of using general-purpose, system-level file systems which may not be optimized for HPC workloads and must be installed by administrators. In contrast, file systems at the user level can be tailored for specific HPC workloads for high performance and can be used by applications without administrator intervention.

Many user-level file system developers have found themselves “having to reinvent the wheel” to implement various optimizations in their file systems. Thus, a main goal of this meeting was to bring together experts in I/O performance, file systems, and storage, and collectively explore the space of current and future problems and solutions for I/O on hierarchical storage systems in order to begin a community effort in enabling user-level file system support for HPC systems. We had a lively week of learning about each other’s approaches as well as unique I/O use cases that can influence the design of a community-driven file and storage system standards.

The agenda for this meeting contained talks from participants on the following high level topics: HPC storage and I/O support today; what do HPC users need for I/O; existing user-level file system efforts; object stores and other alternative storage systems; and components for building user-level file systems. The talks were short and intended to be conversation starters for more in-depth discussions with the whole group. The participants engaged in lengthy discussions on various questions that arose from the talks including: Are we ready to program to a memory hierarchy versus block devices? Are the needs of HPC users reflected in our existing file systems and storage systems? Should we drop or keep POSIX moving forward? What do we mean when we say “user-level file system”? Do we all mean the same thing? How should the IO 500 benchmark be defined so it is fair and useful? and How are stage-in and stage-out actually going to work?

The report for this seminar contains a record of the talks from the participants as well as the resulting discussions. Our hope is that the effort initiated during this seminar will result in long-term collaborations that will benefit the HPC community as a whole.

Challenges and Opportunities of User-Level File Systems for HPC, *Dagstuhl Reports*, Vol. 7, Issue 05, pp. 97–139
Editors: Andrë Brinkmann, Kathryn Mohror, and Weikuan Yu
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Seminar May 14–19, 2017 – http://www.dagstuhl.de/17202
1998 ACM Subject Classification B.4.3 Interconnections, E.2 Data Storage Representations, E.5 Files, H.3.4 Systems and Software
Keywords and phrases High Performance Computing, I/O and Storage, Object Stores, User-level Storage Systems
Digital Object Identifier 10.4230/DagRep.7.5.97

1 Executive Summary

Kathryn Mohror
Andrë Brinkmann
Weikuan Yu

Seminar Overview

The primary goal of this Dagstuhl Seminar was to bring together experts in I/O performance, file systems, and storage, and collectively explore the space of current and future problems and solutions for I/O on hierarchical storage systems. We had a lively week of learning about each other's approaches as well as unique I/O use cases that can influence the design of a community-driven file and storage system standards. We also engaged in several informal, in-depth discussions on questions surrounding how we should best move forward in the I/O and storage community.

A portion of agenda for this meeting was partitioned into sessions containing short talks. The short talk sessions were grouped into high level topic areas: high performance computing and storage systems today; user needs for I/O; user level file system implementations; object stores and alternatives; and file systems building blocks. The intention behind the short talks was to acquaint the attendees with each other's work and to inspire further discussions. Following each talk topic, we had panel-style discussions with the talk speakers serving as the panel. In these panel-style discussions, the audience had the opportunity to ask questions about the speakers' talks as well as note and discuss commonalities and differences across the presentations.

The remainder of the agenda for the meeting was reserved for open discussions with the entire group. The participants engaged in lengthy discussions on various questions that arose from the talks. Additionally, participants were encouraged to propose and vote for discussion topics on a white board. The proposed topics with the most votes were included in the agenda. The in-depth discussion topics included:

- How are stage-in and stage-out operations actually going to work?
- How can we fairly judge the performance of storage systems – IO 500?
- What is a user-level file system? What do we mean when we say that?
- How can we characterize what users need from storage systems?
- Are we ready to program to a memory hierarchy versus block devices?
- and What should we do about POSIX?

The combination of short talks and open discussions resulted in a fruitful meeting. Since the work of the participants was not necessarily familiar to all, the short talks provided a foundation for getting everyone oriented with each other's efforts. Once that was achieved, we were able to productively dive into the informal topic discussions. Overall, several common themes emerged from the talks and discussions. The participants agreed that these themes were important to address to meet the needs of HPC applications on next-generation storage systems. We include these themes in this report in Section 9 to serve as suggestions for further investigations.

Report Organization

Here we present an overview of the topics in this report to guide the reader. Our goal in this report is to capture as much information as possible from the seminar so that those who could not attend can benefit from the talks and discussions.

We detail the short talk sessions in Sections 3-7. First, we provide a summary of the notes from the session note taker and other comments from the talks and panel discussions. Following this summary we provide a listing of each talk in the session and its abstract.

In Sections 8.1-8.6 we give summaries of the informal discussion sessions. The summaries in this case are in outline format in order to capture the conversational and informal nature of the sessions. In many cases, the discussions drew out many interesting questions instead of clear paths forward, so the outline format captured this well.

Following the summaries of the sessions, in Section 9 we conclude with a discussion of recurring themes, including issues for future discussion and work, that occurred during the meeting. We feel that these themes are the true product of this meeting and can serve as a foundation for future meetings or other community efforts.

2 Table of Contents

Executive Summary
Kathryn Mohror, Andrë Brinkmann, and Weikuan Yu 98

Talks Session: High Performance Computing and Storage Systems Today

Some Thoughts on Todays and Future I/O
Thomas Bönisch 107

File Systems for HPC
Mark Parsons 107

HPC File systems Today – What's Working and Opportunities to Improve
Ned Bass 108

Three Musings from Bent at Dagstuhl
John Bent 108

Data Movement Requirements for HPC and Data-Intensive Burst Buffers
Dean Hildebrand 108

Evolving Machine Architectures Are Shifting Our Research Agenda – We Need To Keep Up!
Jay Lofstead 109

A future I/O concept
Franz Josef Pfreundt 109

Talks Session: User Needs for I/O

What do we and the users know about I/O?
Michael Kluge 113

User Workflow Use Cases
David Montoya 113

The SIONlib Multi-file Container Approach for Massively Parallel Task-local I/O
Wolfgang Frings 114

libhio: Optimizing IO on Cray XC Systems With DataWarp
Nathan Hjelm 114

Self-Describing Data
Scott Klasky 114

Talks Session: User Level File System Implementations

Direct-FUSE: A User-level File System with Multiple Backends
Yue Zhu 117

echofs: Enabling Transparent Access to Node-local NVM Burst Buffers for Legacy Applications
Alberto Miranda 117

Exploring User Level Burst Buffer on Public Cloud and HPC
Tianqi Xu 118

MarFS and DeltaFS at LANL: User-Level File Systems Opportunities and Challenges
Brad Settlemyer . . . 118

Gfarm file system meets burst buffer
Osamu Tatebe . . . 119

HPC File System Opportunities with Microkernels (LWK), especially L4-based
Hermann Härtig . . . 119

Talks Session: Object Stores and Alternatives

Proactive Data Containers (PDC) for next generation HPC storage
Suren Byna . . . 121

DAOS: A Scale-Out Object Store for NVRAM and NVMe
Andreas Dilger . . . 121

Are objects the right level of abstraction to enable the convergence between HPC and Big Data at storage level?
Maria S. Perez and Luc Bougë . . . 122

An Exploration into Object Storage
Lance Evans and Raghunath Raja Chandrasekar . . . 122

Talks Session: File Systems Building Blocks

Security Issues in User-Level File Systems
Stergios V. Anastasiadis . . . 128

An effort to systematically understand UFS design requirements
Federico Padua . . . 128

The Case for a Flexible HPC Storage Framework
Michael Kuhn . . . 129

Building blocks for user-level HPC storage services
Philip Carns . . . 129

To FS or not to FS...
Robert B. Ross . . . 130

Discussions

Discussion: How are stage in and stage out operations actually going to work?
Lead: Ned Bass . . . 130

Discussion: How can we fairly judge the performance of storage systems – IO 500?
Lead: John Bent . . . 131

Discussion: What is a user level file system? What do we mean when we say that?
Lead: Philip Carns . . . 132

Discussion: How can we characterize what users need from storage systems? *Lead: Scott Klasky* . . . 133

Discussion: Are we ready to program to a memory hierarchy versus block devices?
Lead: Jay Lofstead . . . 133

Discussion: What should we do about POSIX? *Leads: Rob Ross and Andreas Dilger* 135

Recurring Themes and Future Work from Seminar

Recurring Themes 136

Future Work 138

Participants 139

3 Talks Session: High Performance Computing and Storage Systems Today

HPC is undergoing several transformations at the same time, where many of them include storage systems and technologies as a driving component:

- The growing gap between the performance of magnetic disks and processors drove the introduction of Solid State Drives (SSDs) as an additional layer in the storage hierarchy. SSDs are more expensive for the same storage capacity than traditional magnetic disk drives, but they are more cost efficient with respect to small I/O performance. Back end parallel file systems of HPC clusters with capacities in the range of tens to hundreds of PBytes are therefore still mostly based on magnetic disks, while SSDs are used as burst buffers to cache data, which can be implemented as either node-local storage or as shared storage external to compute nodes.
- The workload on many HPC systems is changing from pure number crunching to also include data intensive computing, which uses data parallelism to process data sets with a size of many PBytes per application run. Principles of data locality and data caching borrowed from big data approaches help to improve overall application performance, making burst buffers with low access latencies especially attractive.

Therefore, the first session of this Dagstuhl seminar focused on the big picture of HPC architectures and the current (and possible future) usage of HPC storage systems.

Thomas Böhnisch started with a presentation of the storage system at the High-Performance Computing Center HLRS in Stuttgart. The center focuses on engineering and industry applications and 80% of the applications periodically output data, while this data is not necessarily intended as defensive checkpoint data. In his talk, **Some Thoughts on Todays and Future I/O** Thomas described use cases, where applications spend more than 1.5 days of 2 days runtime in I/O even after code optimization, which had been able to increase file system bandwidth usage from 1 GByte/s to 7.5 GByte/s. The bandwidth usage nevertheless is only 1/10-th of the backend storage bandwidth available in Stuttgart, where the storage systems delivers up to 75 GByte/s. The main reason seems to be that many applications developers only optimize on one aspect of the code, which rarely is I/O.

Applications seem to become, according to Thomas, a black box for the scientific users, who have no knowledge of the I/O performed by the application and even much less of optimal file system settings. Thomas presented burst buffers as a possible way to overcome challenging and bursty I/O patterns, but also mentioned that they have to be seamlessly integrated into the HPC environment, as users want to perform their science and as they are not interested in I/O. Furthermore, he identified storage class memory as a possible game changer, which will impact the working methods in HPC, while there is quite some research and development necessary to efficiently use it.

Mark Parson from the Edinburgh Parallel Computing Centre (EPCC) started his talk **File Systems for HPC** with the observation that the complexity of storage systems has greatly increased in the past decade. Most HPC systems now have multiple storage systems and multiple file systems per storage system, while GPFS and Lustre dominate as backend parallel file systems. Resiliency stays an important aspect of storage management and most HPC centers have at least one storage crisis a year.

The dominant usage scenario seen in Edinburgh is that most users use the POSIX file system interface and do not include parallel IO libraries. Layered I/O (NetCDF over HDF5 over MPI-IO) approaches on the other side fight with each other by providing optimizations at each layer.

Buying storage is, according to Mark, still terribly complicated and confusing. The bandwidth requirements are mostly dominated by the start and end-phases of jobs as well as by checkpointing, which are very read-write-intensive, and which therefore waste investment. Furthermore, performance degrades over time (e.g., due to file system aging effects) and scientific users often *miss-use* parallel file systems. Examples from genome processing include that one user moved around 400 TB through storage every week, while others created more than 240 million files in a directory within one year, stressing both Lustre and GPFS. It would be therefore beneficial to have parallel I/O benchmarks, which can be used for procurement purposes, including different HPC usage scenarios.

Mark's talk agrees with Thomas in that a new hierarchy of next generation NVRAM technologies will profoundly change memory and storage hierarchies, while HPC systems and data intensive computing systems will merge. Profound changes will therefore come to all data centers, while HPC centers still need to develop software – OS and application – to support their use. This change of HPC storage has to be seen in the context that parallelism beyond 100 million threads demands new I/O approaches, where reading and writing data to parallel file systems still seems a major bottleneck and data movement and processing on capability systems has to be re-thought. Mark sees a demand for truly parallel file systems with reproducible performance, while current technologies simply will not scale, leading to large jobs spending hours reading data and writing results.

The talk **HPC Filesystems Today – What's Working and Opportunities to Improve** from Ned Bass is based on the experiences gained running seven Lustre file systems being in production at Lawrence Livermore National Laboratory, hosting more than 50 PByte on the open side. Ned started with a summary of well-working aspects, including the Open Source development of Lustre and ZFS (even while pushing changes up-stream is difficult). Stability has furthermore improved according to him a lot compared with the past and it is difficult for him to remember the last time when LLNL has lost user data. Furthermore, it is in many cases not the fault of the file system if something goes wrong, but even in such cases users often perceive the file system as a root cause for a problem. Another advantage of today's storage systems is that the POSIX programming is well understood and is portable from a laptop up to a supercomputer. Relaxing POSIX would therefore also mean that it becomes necessary to build a compatibility layer so that people can still perform application development on their laptops.

A further observation from Ned is that well-formed I/O also performs well. Nevertheless, the definition of *well-formed* depends on the file system in use and is not even portable between Lustre and GPFS. Additional challenges are, according to Ned, that the storage hierarchy is not transparent to the end users and that it is a heavy burden for them to understand where data lives and where it *should* live. Ned sees inflexible semantics, where the system decides where the data stays, as one of the obstacles and wants a clean interface for moving data between the storage hierarchies.

He agrees with the former speakers that users do not care about I/O and assume the technical staff to make it work and fast. Users spend very little time on I/O optimization, as long as it stays below 10% of the runtime they do not care. Managing data on the other side (including data visibility and debugging) requires a high degree of expertise from developers and system administrators, who could benefit from an API covering the complete storage hierarchy. Deciding on such an API is on the other side a big investment, as it can incur high technical debts if it is not well thought through.

The influence of SSDs and burst buffers on the storage architecture of HPC centers has been the core of John Bent's talk **Three Musings from Bent at Dagstuhl**. The talk describes the change of the storage architecture, starting with the two-layer architecture consisting of parallel file systems and tape archives, which has been dominant from the early 2000's until 2015. Both, aligned and unaligned storage accesses hit the parallel file system, while I/O to tape has been typically well-formed (based on any definition).

The price advantage of SSDs concerning I/O rate and bandwidth led to the introduction of burst buffers, which are able to absorb the unaligned IO traffic, while changing this traffic into well-aligned accesses to the parallel file system. The next layer, which is already becoming part of many HPC installations, is an object storage layer for data which is neither hot (and therefore placed on the parallel file system or the burst buffer) nor cold (and therefore being stored on tape). This lukewarm storage occurs in the context of simulation campaigns, where huge data volumes are analyzed over time spans of many months.

Object storage has become prominent in the context of Cloud computing and has the advantage of nearly unlimited scalability, while offering the relatively low latencies and sequential access bandwidth of (highly-parallel) magnetic disks. A major additional advantage of object storage is furthermore its manageability, which, according to John, helps to decrease overall cost even below tape costs.

These four layers are (from most attendees' perspectives) too many layers for users. John's proposal is therefore to radically change the storage architecture and to drop the backend parallel file system and the tape archive and to keep hot data in burst buffers, while storing all cold to lukewarm data in object storage. John closes his presentation with a comparison of different possible locations of burst buffers, including node local storage as private burst buffers (e.g., in Cray/Intel Aurora at Argonne), as node-external or shared burst buffers (e.g., deployed in Cray Trinity at LANL) or embedded burst buffers (e.g., being part of Seagate Nytro NXD).

The talk **Data Movement Requirements for HPC and Data-Intensive Burst Buffers** from Dean Hildebrand focuses on the requirements to integrate burst buffers into the HPC environment. He argues similar to the previous speakers that burst buffers (and therefore SSDs) have not mainly been integrated based on their superior performance, but mostly based on their lower costs per (unaligned) IOPS. Dean therefore discusses workloads and memory access models, which are suited for burst buffers and asks for the necessary namespace integration approach, the durability, capacity, and availability requirements and the best API to use burst buffers, whether it might be POSIX, key value or object.

Transferring data in and out of burst buffers is seen as an unsolved problem by Dean. Example applications from the domain of deep learning, like Caffe, Attila, or TensorFlow are putting high pressure on the storage backend, where input data sets range from ten GByte up to tens of TByte. The entire input data sets are re-analyzed dozens of times and even intermediate writes can easily grow to tens of TBytes. Furthermore, it is often not clear which data is required where, so that all data is copied to all nodes.

The talk therefore identifies several requirements to efficiently integrate burst buffers into HPC designs, especially in the context of new applications. It is obviously important to reduce the bandwidth between burst buffers and the durable tier, while it can be very beneficial to include application specific direct data placement. Burst buffers can furthermore only be efficiently used if the network bandwidth scales with the capacity of the burst buffer. Unsolved challenges also include efficient scheduling strategies (and to couple them with HPC batch environments), which help to identify for each data set when it has to be transferred where. Especially node-external data sets also impose new security challenges and demand for multi-tenancy features.

In summary, Dean sees burst buffers as caches, not primary data stores, which are useless without data, while efficiently transferring data in and out of burst buffers seems to be still an unsolved challenge.

The talks discussed so far focused on technology transitions, which are already visible in today's HPC environments. Jay Lofstead's talk **Evolving Machine Architectures Are Shifting Our Research Agenda – We Need To Keep Up!** discusses the impact of new memory and storage technologies based on byte addressable storage class memory (SCM), which builds another layer into the memory/storage hierarchy. SCM is able to not only blur the dividing line between memory and storage, but to obliterate it.

A challenge seen by Jay is that the HPC community has not yet adequately solved the problems inherent to the architectures being deployed today, not to mention those of the future, where memory becomes persistent and networking becomes part of the memory hierarchy instead of just the storage hierarchy.

Fundamental questions during the transition to SCM include the interface to SCM. Can we keep files as an abstraction if the POSIX semantics cannot be kept without inducing major performance losses? How do we identify a collection of bytes in the absence of files? Do we still call it storage if we use CPU-level get/put operations? Additional challenges during the transition include keeping consistency, coherence, and security, which are currently provided by file systems.

Jay describes four transition phases, where phase one already started in the late 1990's, when extra compute nodes have been used for their memory as data staging areas, building the origin of *burst buffers*. Flash as burst buffers in or near the IO path has been the starting point for phase 2, including rudimentary job scheduler support for data pre-staging, and data draining.

Phase 2 is still on-going and explores additional use cases for Flash, e.g. as part of the compute nodes. This new architecture might lead to interferences, as accesses impact the network, memory or disk bus of the local node. Phase 3 will provide additional node local storage, e.g. as persistent 3D Xpoint, while nodes also gain on package high-bandwidth memory (HBM). Challenges in phase 2 and 3 are the interconnect speed, erase cycles in Flash and 3D Xpoint, maintaining coherency and consistency for multi-user support and globally shared spaces.

Phase 4 of the architectural transition will then provide a memory-centric design with network storage (on switches), DRAM (potentially in the same address space). The line between memory and storage will therefore be all but gone. Typical use cases will be coherent virtual fat nodes operating on 10s of TBytes, the introduction of persistent storage near/fast enough to "swap" to, while hopefully providing new and easier programming models to access data.

A possible dividing line between memory and storage might be (according to Jay) that (today) things placed in memory have external metadata, which is generally expressed in the program code. Memory therefore has a compact representation, which is optimized for the interaction with processors. Things placed in storage are (today) wrapped in metadata to make them easily usable by other applications. File formats make it possible to read in simulation output into visualization tool, even in case of a different endianness.

Franz-Josef Pfreundt tries to tackle the new challenges from two different perspectives. His talk **Thoughts about the future of IO** presents BeeOND (a sibling of BeeGFS) as a burst buffer file system to efficiently use node internal SSDs. BeeOND can build a temporary file system across nodes (on demand per user), where every compute node can become a metadata server. Possible approaches, e.g., to deep learning (DL) I/O challenges are to combine all files in one large binary with fixed offsets.

Data management is seen by Franz-Josef as one key problem in parallel computing, as it is too complicated for application developers. He therefore also proposes new communication interfaces (in his case the GPI-2 Global Address Space Communication Interface) with explicit one-sided communication with notification.

GPI provides a standardized API (GASPI) and hides latencies by asynchronous one sided RDMA communication. Memory virtualization is provided by GPI Spaces, which can keep data without a running application being responsible for the data. Data exchange between applications and the virtual memory is performed through shared memory segments, while data transfer between nodes is provided by GPI. GPI Space is currently becoming a complete distributed runtime system, which provides failure tolerance, JIT compilation and execution of the underlying Petri nets, co-scheduling of multi-node tasks and preemptive scheduling of data transfers.

The presentation included two use-cases, where the combination of GPI-Space and domain knowledge helped to overcome application problems. SPLOTCH, a MPI visualization program from the astrophysics domain, has been rewritten in GPI-Space in three weeks, while being able to improve the application runtime by a factor of more than three in the worst case. A Deep Learning on Demand Architecture has been presented as a second use case, which uses Amazon's spot market to provide automatic meta-parameter search and which offers automatic data-management, while still supporting the original DL model descriptors, e.g. from Caffe or Tensor Flow. Future work includes an architecture that offers virtual memory directories and caches based on a client-server architecture, while automating the transfer between storage and (persistent) memory.

3.1 Some Thoughts on Todays and Future I/O

Thomas Bönisch (HLRS – Stuttgart, DE)

The presentation shows the three main issues of user applications with a current file system implementation at the Hazel Hen HPC system in Stuttgart. The first issue is that applications very often cannot make use of the file system performance. We discuss the reasons for that depending on the user group. The second issue is that I/O is in bursts and third, several applications tend to generate millions of files. In addition to these issues, we discuss a new approach for data management and what this could mean for user applications.

3.2 File Systems for HPC

Mark Parsons (University of Edinburgh, GB)

From a data centre point of view the quantity and complexity of data we are being asked to manage is growing all the time. Coupled to this are complex requirements around security – particularly related to authorisation, accounting and auditing – who uses what data, when, and why. In the High Performance Computing and Data Analytics world, new memory technologies are on the immediate horizon that will markedly change how use and manage data in both the research and commercial worlds. We need to think carefully about what these competing demands mean for file systems and how we can provide rich environments for data manipulation and its transformation into knowledge in whatever sector it is being used.

3.3 HPC File systems Today – What's Working and Opportunities to Improve

Ned Bass (LLNL – Livermore, US)

In this talk I provide an overview of the HPC storage environment at LLNL, and I discuss successes and challenges with the current state of the art. Aspects that are working well include stability, data integrity, security, scalability for petascale class systems, a well understood portable programming model, and thriving open source collaborations. Some major challenges also exist in HPC file systems today. Users are exposed to internal implementation details of the storage hierarchy which would be better hidden behind intuitive abstractions. Options to tailor API semantics and consistency models to the needs of the application are limited. Users bear a heavy burden to curate their data and manage its movement throughout the storage hierarchy. A well-designed interface would allow users to provide hints about the expected life span and access time requirements of their data and the system would do the right thing on their behalf. Finally, due to their complexity and aggressive pace of development, current HPC file systems carry high total cost of ownership and significant technical debt. Now is the time to lay the groundwork for the maintainability of future systems through careful design and thorough documentation. As we move toward new storage paradigms to address current challenges we should take care not to lose ground in the areas that are working well today.

3.4 Three Musings from Bent at Dagstuhl

John Bent (Seagate Government Solutions – Herndon, US)

One, the grind-crunch model of computational simulation and what would you do with one more dollar. Two, from 2 to 4 and back again; the evolution of HPC storage architectures. Three, to share or not to share; a comparison of burst buffer architectures.

3.5 Data Movement Requirements for HPC and Data-Intensive Burst Buffers

Dean Hildebrand (IBM Almaden Center – San Jose, US)

Over the next several years, all applications—regardless of industry—will begin to expect the bandwidth and latency of fast storage devices such as NVMe and 3D XPoint. A discrepancy therefore exists between the high-performance expectations of users vs. what their budgets can contain. This means that dataset movement between slower durable storage and fast storage will become an integral part of the storage hierarchy on-premise and in the cloud. This is further complicated by the diverse set of storage software and their APIs that exist at both the fast and durable tiers (e.g., NAS, object, block, parallel file systems, tape, optical), most of which struggle to communicate and share data. In my talk, I will discuss the growing relevance of "bust-buffer'-like data movement requirements across a variety of industries and their emerging requirements.

3.6 Evolving Machine Architectures Are Shifting Our Research Agenda – We Need To Keep Up!

Jay Lofstead (Sandia National Labs – Albuquerque, US)

While IO/storage research has slowed in recent years, the bits that still come out tend to be clustered in incremental improvements to two-phase collective IO, isolated traditional storage aspects such as metadata improvements, or trying to understand the chaotic IO patterns in an attempt to predict and/or schedule IO more efficiently. While there is some value in these research areas, they are missing the real problems we will face soon – changing underlying storage/memory technologies and machine architectures. With SSDs creeping into HPC platforms and NVM over fabric being a near term product, we in the IO community need to focus more efforts to address these challenges or risk losing relevance to the programming models community. This talk will explore why the current topics outlined above are less useful or even harmful to continue to investigate while laying out a case for why we need a serious shift in focus to stay relevant.

3.7 A future I/O concept

Franz Josef Pfreundt (Fraunhofer ITWM – Kaiserslautern, DE)

With the appearance of large memory machines and the future dominance of fast storage class memory we have to rethink parallel I/O and parallel programming. At Fraunhofer we work on two software systems that are intertwined with respect to this topic. GPI-Space is a system to automate parallel programming in using a global virtual memory space and petri nets to express parallelism. Today we can map already the virtual memory space into our parallel file system to extend the size and add failure tolerance. BeeGFS, the parallel file system that we develop, is a state of the art POSIX compliant parallel meta file system. How we can merge the two systems is our topic for the future. At the end, the goal must be to manage data transport from disks to HBM fully automatic.

4 Talks Session: User Needs for I/O

Technological progress can only benefit us if it is in line with our needs. New device technologies and faster supercomputers are therefore only useful if they enable users (here in the form of scientists) to perform better research. This includes several different dimensions.

First of all, science has to demand for faster supercomputers to justify growing investments. The widening HPC user community in research as well as the broad application of HPC in industry seems to be a good indication for this demand. Second, it is important that the new capabilities are provided to the scientist in a way that they can be easily integrated into their scientific workflows, so that scientists do not have to become HPC experts, which would distract them from doing the research. Third, applied HPC research should focus on the current pain points of researchers, optimizing scalability challenges.

The presentation **What do we and the users know about I/O?** from Michael Kluge started with an overview about I/O usage data which is typically collected and analyzed in HPC data centers. The reason for this data collection is that I/O is a shared resource, where individual users can abuse the file system in a way that interferes with concurrently running jobs of other users. This interference can become in practice a denial of service attack, as HPC users have often not been taught to program in an I/O friendly manner, while there mental model of I/O is not compatible with the real HPC (storage) hardware.

The positive news is that there are today plenty of data collection tools available, which help to get an overview about, e.g., how much data has been read or written from which user at a specific bandwidth. There are also many performance counters integrated, helping to also get detailed information about additional metadata usage. This profile data helps system administrators to get a good overview about quantity and quality of storage usage, but it does not immediately help us to learn about which new interfaces are needed and how we can modify current applications to become more I/O-friendly.

A hypothesis within Michael's talk was that we do not only need new file system interfaces, but that many new storage interfaces with different semantic variations of the POSIX interface are required. New semantics can help to find a sweet-spot between avoiding performance bottlenecks (which are induced by today's file system implementations) and changing as little as possible for users. The main questions seem to be how much the file system interface has to be changed for specific applications and whether application experts, users, and system administrators are able to give convincing answers to this question.

It seems obvious that the answers to these questions are not easy, as changing the API requires knowledge about what semantics an application relies on when using a storage interface. This does not only includes the file system itself, but also I/O libraries and abstraction layers like burst buffers, I/O forwarding nodes, or automated workflow engines. It is very likely in this context that the knowledge about I/O usage decreases with an increasing depth of the I/O stack.

Michael proposes two possible approaches (as well as combinations). The formal approach requires the developer to describe the kind of features needed, while I/O libraries and file systems describe the features which they provide. Both lists can be matched with each other, offering to the application the storage interface with the minimal set of features required (and therefore also hopefully with high performance). The technical approach tackles the challenge per use case by system wide I/O tracing and analysis. This tracing can then hopefully reveal all dependencies, but would be limited to this specific execution.

Dave Montaya highlights in his talk **User Workflow Use Cases** a trend seen in many HPC data centers: physical insights are not gained by a single application run, but by an ensemble of applications working together within a workflow. Running large suites of simulations is according to Dave already routinely done at the Los Alamos National Laboratory (LANL) to study the effects of (usually small) changes to problem geometry and/or materials. Furthermore, workflows help to continually assess the viability of new physics models, meshing strategies or code settings that are not related to physical changes to a problem while they enable rapid simulation turnaround, repeatability, and consistency.

The data infrastructure is according to Dave one of the pain points, as it is continuously evolving. The well-understood hierarchy of memory, parallel file systems, and archive is already including additional storage layers, like burst buffers and campaign storage, while the boundaries between these layers will blur with the introduction of byte-addressable storage class memory. These changes in the infrastructure have also to be addressed by changing the workflows underlying the simulation runs.

Dave also sees challenges concerning the large sets of input data with complex dependencies. This input data includes the syntax of the physics code, the common model data, as well as suite specific data, which overrides the common model and gives explicit meshing instructions and provides execution strategies, while the shot specific data includes additional overrides for the specific workflow run.

These challenges have to be seen in the context of the size of large scale capability workflows. Dave explains, based on the example of the Hydra UQ workflow to support ignition experiments at the National Ignition Facility (NIF) in Livermore that each study run can include up to 60,000 simulation runs within each single allocation, generating over one billion synthetic x-ray images over a runtime of 8 weeks. Each of these simulation runs generates new files every few minutes, leading to more than 5 PBytes of intermediate files. Even if the average bandwidth of 40 GByte/s does not saturate the backend Lustre file system, it would have already been very useful to be able to integrate burst buffers running at higher peak bandwidth.

The talk **The SIONlib Multi-file Container Approach for Massively Parallel Task-local I/O** by Wolfgang Frings moves from the problem definition to providing solutions for a specific challenge. Many applications running on JUQUEEN in Jülich (and on many other HPC centers) generate task local checkpoint files, while I/O is handled by separate I/O nodes in each rack. Generating a checkpoint from applications running on up to 458,752 cores can therefore easily flood the I/O nodes as well as the underlying parallel file system. Wolfgang presented graphs showing that generating one file from each JUQUEEN core can already take more than four minutes, while assuming 64 tasks on each node easily leads to file `create()`-times of more than 10 minutes.

The main approach of SIONlib is to provide an interposing layer, which transparently translates the $1:1$ relationships between clients and files into a $n:1$ relationship, so that n clients share 1 file. The number of new files can therefore be reduced to a few, significantly decreasing the `create()` time. SIONlib shows on the other hand a tradeoff concerning the number of files, which are shared by many clients, and the write bandwidth accessing these files. Wolfgang showed that the `create()`-times can be reduced to a few milliseconds if all threads share a single file, while at the same time the bandwidth to this shared file significantly decreases with the number of accessing tasks.

SIONlib therefore introduced a multi-file approach, which maps clients to I/O nodes, so that no I/O node gets overwhelmed by accesses of too many clients. Applications with an 1:1 access pattern like checkpointing still work by simply linking the SIONlib library achieving up to 67% of the peak bandwidth for 1.8 million tasks. SIONlib is already integrated into many HPC applications like multigrid solvers (DUNE-ISTL), fusion simulation (ITM), or applications to simulate the human brain (NEST).

The new Buddy CP approach also integrates local storage within cluster nodes to further increase possible bandwidth utilization. Data is first written to the local storage of the node, where a process is running, while it can be later (transparently) redistributed to a buddy node within the cluster or to the global file system.

Nathan Hjelm's talk **libhio: Optimizing IO on Cray XC Systems With DataWarp** tackles a similar challenge from a different perspective: How do we use new technology if it is provided to us? Cray's DataWarp technology is a burst buffer implementation, which has been initially developed for the leadership computers Trinity at Los Alamos National Laboratory and Cori at NERSC. The DataWarp implementation at LANL consists of service nodes including one Xeon E5 processor and two 3.2 PCIe SSDs each, which are directly connected to the Cray Aries network. The software stack includes an API with functions to initiate stage in / stage out and to query stage state. DataWarp can be configured in multiple modes using the workload manger, e.g., to stripe data over multiple nodes, to use it as a cache or in shared / private mode.

The Hierarchical IO Library (HIO) has been developed at LANL to facilitate the use of burst buffer technologies and specially to decouple application development from today's storage implementations, so that they can still efficiently run on next generation hardware. The development has been user-focused, so that it can be easily incorporated into existing applications and help to improve checkpoint performance for $n : 1$, $n : n$, and $n : m$ IO patterns.

HIO is centered around an abstract name space and each HIO *file* is a *dataset*. HIO datasets support named binary elements and a shared or unique-per-rank element offset space, where an offset space is defined per named element. The internal on-disk structure is not yet specified to preserve flexibility for future optimization. The currently supported "File Per Node"-mode has been originally targeted for Lustre and is intended to reduce the metadata load compared to traditional $n : n$ approaches, while leading to less file contention than the $n : 1$ approach, leading to similar optimizations like presented for SIONlib. Cross-process data lookup is supported via MPI-3 Remote Memory Access (RMA).

Scott Klasky's talk on **Self Describing Data** abstracted from today's user requirements and provided an outlook on how data can be better (re-)used. The talk started (after a short outlook on ADIOS 2.0) with the observation that raw data is of little use if it is not accompanied with an appropriate description. The costs for recovering information from such datasets typically increases linearly with the size of the dataset. Self-describing datasets should simplify the access to information by annotating the data with easily and fast accessible metadata.

An important challenge is to find the extra annotations to make data more valuable without adding so much information that it greatly affects performance. Self-describing data is thereby nothing completely new, examples of self-describing files include MS Word, JPEG, HDF5, or ADIOS. ADIOS, e.g., includes a self-describing BP file format, which enables to read arbitrary subarrays of variables, so that variables being written out from n processor can be read in by an arbitrary number of processors.

Self-describing data is especially useful for queries, as searching through the data can be largely simplified by its self-descriptive nature. Data analysis for decision-making can therefore be performed on a "need to know" information basis. Metadata can include information on min/max-values of temperatures, power consumptions, or humidity for certain data blocks, so that blocks with important data can be easily derived and further analyzed. It is also in principle possible to put time constraints on queries, e.g., to get the best possible information which can be collected in one hour.

Potential performance problems using self-describing files include that it can become expensive to capture global metadata when the problem size grows. ADIOS partially solves the problem by maintaining redundant metadata. The redundancy mechanisms allow ADIOS to regenerate metadata in case of failures and in case of performance issues when writing. Metadata can furthermore kept locally, while a (virtual) global metadata file can be generated for online operations.

(Virtually) Global metadata makes it possible to first read through the metadata and then to only move data needed based on metadata operations, which is important for scientific workflows including multiple sites. The proposed metadata analysis is in this case cheap compared to the costs incurred by blindly moving huge datasets. Using self-describing data for staging can enhance data services and communication among applications providing an intermediate common area (staging) that reduces file system access costs, while it can include plugins to perform analytics and visualization, data reduction, or data transport.

Coupling tools like ADIOS or HDF5 enables application developers to present their data based on a common schema, so that tools in a workflow can cooperate with each other without providing data transformations for every new coupling between applications. Refactoring data to fit into schemas also enables applications not only to read data which is important, but also to perform on the fly transformations, e.g., by reducing the accuracy of certain variables, as this does not only decreases transfer, but also later processing costs.

4.1 What do we and the users know about I/O?

Michael Kluge (TU Dresden, DE)

Research file systems relax POSIX semantics in different ways. The questions is, whether there can be a common description about what semantics a certain file system provides. My talk has an proposal about what kind of information is needed to make proper decisions about whether or not a research file system fits to a given use case.

4.2 User Workflow Use Cases

David Montoya (Los Alamos National Lab., US)

Understanding simulation study workflows has become more important as HPC architectures and computational scales increase. There is a need to communicate application patterns and behaviors to vendors and performance insight to application developers. We have started a methodology and tools to describe and collect metrics to support this need. Additionally we describe use cases that show how workflows, which include pipelines of application runs are evolving and changing to adapt to architecture changes. The data management and provenance of these workflows are creating the need for a distributed view of data interdependencies.

4.3 The SIONlib Multi-file Container Approach for Massively Parallel Task-local I/O

Wolfgang Frings (Jülich Supercomputing Centre, DE)

The talk will focus on the multi-file container approach of SIONlib, a parallel I/O library, which addresses the problem of large number of task-local files by transparently mapping the task-local files onto a small number of physical files. Among other optimizations SIONlib uses a virtual shared container which is divided into a small number of physical files to represent a user based file system for parallel task-local I/O on application layer. Using multiple shared files instead of one big file allows to benefit from hierarchical storage systems, for example, by managing these files on node-local storage. In addition, recent developments within the EU-project DEEP-ER will be shown, where a combination of SCR and SIONlib is used to implement efficient multi-level checkpointing based on shared file containers. Within this project SIONlib was extended with a checkpointing option, based on the multi-file approach, which allows to store additional copies on buddy-nodes or to store parity information on network attached memory.

4.4 libhio: Optimizing IO on Cray XC Systems With DataWarp

Nathan Hjelm (Los Alamos National Laboratory, US)

High performance systems are rapidly increasing in size and complexity. To keep up with the Input/Output (IO) demands of High Performance Computing (HPC) applications and to provide improved functionality, performance and cost, IO subsystems are also increasing in complexity. To help applications to utilize and exploit increased functionality and improved performance in this more complex environment, we developed a new Hierarchical IO (HIO) library: libhio. In this paper we present the motivation behind the development, the design, and features of libhio. We detail optimizations made in libhio to support checkpoint/restart IO workloads on the Trinity supercomputer, a Cray XC-40 at Los Alamos National Lab. We compare the large scale whole file read/write performance or IOR when using libhio against using POSIX when writing to Cray DataWarp.

4.5 Self-Describing Data

Scott Klasky (Oak Ridge National Laboratory, US)

There is a critical national need to enhance economic development using high-performance computing. The exascale challenge will transition HPC science to a "predictive" science which means we must be able to extract and preserve sufficiently accurate output from simulations. Exascale systems are projected to provide a far greater increase in computational speed than in I/O bandwidth, compared with petascale systems. This creates a time-critical need to optimize I/O to make sure that essential data generated by science applications can be written and read from the storage system. The data generated by these applications are extremely valuable, since the cost of running at scale on the average system on DOE cost over $200K USD per day (today), and is expected to become even more expense. Therefore, we argue that the data should be annotated in a self-describing storage format, such as ADIOS, HDF5, netcdf, etc. There are many challenges to do this at scale, and make this ubiquitous among HPC applications, but generally these lead us to understand how to make this easy to use and easy to achieve extreme portable performance for "most" of the I/O patterns from large scale scientific applications. Furthermore, since many of the newer exascale requirements are pushing us into the land of on-line analysis, we argue that the data formats used in these streaming/staging techniques should also use this self-describing data streams to allow a new data software-eco-system to develop and be nurtured through our efforts.

5 Talks Session: User Level File System Implementations

The third session of the Dagstuhl seminar focused on the development of user-level file systems to meet various needs of HPC applications and systems. Yue Zhu a second-year doctoral graduate student from Florida State University started this session with her talk on **Direct-FUSE: A User-level File System with Multiple Backends**. Her talk begins with a comparison of user-level and kernel-level file systems on their development complexity, reliability and portability. Yue enumerated a number of file systems that are implemented as user-level libraries to avoid the involvement of FUSE kernel module, but argued that an application may face problems when having to work with multiple backend file systems and their distinct file paths and descriptors. She proposes Direct-FUSE to provide multiple backend services for one application without going through the FUSE kernel. While FUSE alleviates the difficulty of developing a kernel-based file system, she recognizes the need of addressing the significant overheads for FUSE I/O function calls. Through a detailed analysis on the costs from context switches, data movements and metadata operations, Yue elaborated her approaches to alleviating the costs and achieving an efficient implementation of Direct-FUSE. Her results demonstrate an effective prototype of Direct-FUSE that supports multiple backend file systems.

Alberto Miranda from the Barcelona Supercomputing Center of Spain then presented his talk on **echoFS: Enabling Transparent Access to Node-local NVM Burst Buffers for Legacy Applications**. He reiterated the I/O challenge faced by petascale and future exascale systems and compared the two options of introducing burst buffers into HPC systems for data-intensive simulation and analysis. Alberto considers that node-local NVM-based storage to compute nodes provides denser burst buffers and offers an opportunity to construct temporary file systems with ad-hoc I/O optimizations for specific batch jobs. Through a research project NextGenIO that is funded by the European Union, Alberto and his team introduce a new user-level file system called echoFS that aims to aggregate the NVM burst buffers available to compute nodes into a collaborative burst buffer and then coordinate with the job scheduler to preload a job's data dependencies for first run-time I/O and simulation performance. While echoFS can benefit from explicit user inputs, echoFS can glean implicit hints from the SLURM batch submission scripts without placing a burden on the application users. In addition, echoFS a a pseudo-randomized striping strategy that aims at distributing data across nodes for elastic load balance, and scalable aggregated I/O, within the same job or across different jobs. Furthermore, a POSIX interface provided from echoFS allows easy access to burst buffers available from any compute node (e.g. SSDs or NVM) for legacy applications. While the echoFS project is still ongoing research, it has demonstrated a number of attractive features that a user-level file system can achieve for ad-hoc application users.

This session continues with more talks on the implementation of user-level file systems. Tianqi Xu a doctoral student from Tokyo Institute of Technology in Japan presented his talk on **Exploring User Level Burst Buffer on Public Cloud and HPC**. To exploit burst buffers for fast I/O required by large-scale applications, Tianqi explored the concept of user-level file systems in two distinct computing environments: public clouds and HPC centers. For the first environment, Tianqi and his colleagues introduce CloudBB as a user-level and on-demand file system to leverage burst buffers in the cloud, tapping virtually unlimited resources therein. Tianqi positioned the burst buffers managed by CloudBB as a new storage tier shared by compute nodes in the cloud. These burst buffers, in combination with the shared cloud storage, form an on-demand two-level hierarchical storage system managed by CloudBB. Together, this new user-level file system enables scalable I/O with multiple metadata servers, support fault resilience through file replication and recovery techniques. Besides its support to burst buffers in the cloud, Tianqi extends CloudBB with additional capabilities to exploit burst buffers for applications in the HPC centers. His work is then manifested into a new implementation called HuronFS. By building dedicated burst buffer on-demand, HuronFS can accelerate applications and avoid I/O contention by leveraging low-latency and high-bandwidth interconnects (InfiniBand) commonly available from HPC centers. It has been shown to work very efficiently on Tokyo Tech's supercomputer, TSUBAME2.5.

The potential of user-level file systems is further demonstrated with two file system implementations MarFS and DeltaFS developed by a team from the Los Alamos National Laboratory in the U.S. Further in this session, Brad Settlemyer from that team covered these two file systems with his talk on **MarFS and DeltaFS at LANL: User-Level File Systems Opportunities and Challenges**. Aiming for great simplicity, MarFS leverages a tool called PFTool that provides parallel file transfer agents (FTAs) to distribute storage objects across a scalable long-term storage system. All I/O operations have to go through the FTA front engine, where the small files will be packed and large files will be sharded. The resulting MarFS is capable of storing extremely large scientific data sets with really scalable metadata performance. Brad argues that MarFS is able to combine a simple POSIX consistency model while enable simple capacity addition over time. Furthermore, Brad's team has designed another user-level file system called DeltaFS that decouples application metadata load from the underlying HPC platform storage system, thereby enabling an extremely scalable metadata plane for VPIC (Vector Particle in Cell) applications. Brad believes that both systems allow LANL to procure storage systems with great flexibility while supporting ultrascale scientific simulation workloads.

Besides enabling new file systems for burst buffers, the power of user-level file system has also been employed for existing file systems to leverage burst buffers. Osamu Tatebe from the University of Tsukuba in Japan presents a case study along this line through his talk on **Gfarm file system meets burst buffers**. Like Google FS and HDFS, Osamu and his team have developed their own high-performance file system called Gfarm to exploit local access performance and avoid network data transfer. His team sees a new opportunity to leverage node local NVMe SSD and NVRAM. They contend that the importance of durability constraints can be downgraded while striving for storage performance, especially for temporary data. To this end, Osamu has modified the conventional I/O calls such as pread and pwrite to leverage fast RDMA operations and dynamic memory registrations, while ditching the durability for temporary data. The experimental results show that this extension can enable effective integration of burst buffers in Gfarm.

Finally, Prof. Hermann Härtig from TU Dresden in Germany presented his talk on **HPC File System Opportunities with Microkernels (LWK), especially L4-based**. Hermann argued that micro kernels have a lot to offer for the construction of user-level file systems. With a review of HPC operating system research, he pointed out that a variant of systems has become popular again after a period of hibernation, e.g., by combining a micro kernel with monolithic OS while splitting the application functions across the two. He then contends that a number of important but interesting questions ought to be addressed for the resulting software ecosystem to function efficiently. A lot of these questions rest with how the system calls can be implemented across the split between the micro kernel and OS. There are ample opportunities for OS experts as well as researchers on user-level file system implementations.

5.1 Direct-FUSE: A User-level File System with Multiple Backends

Yue Zhu (Florida State University – Tallahassee, US)

Although the FUSE file system alleviates the difficulty of developing a kernel-based file system, it introduced significant amount of additional overheads for I/O function calls. Some file systems are leveraged as libraries to avoid the involvement of FUSE kernel module. However, they may meet problems when dealing with multiple backends and distinct file paths and file descriptors for different backends. We propose Direct-FUSE to provide multiple backend services for one application without going through the FUSE kernel.

5.2 echofs: Enabling Transparent Access to Node-local NVM Burst Buffers for Legacy Applications

Alberto Miranda (Barcelona Supercomputing Center, ES)

The current growth in data-intensive scientific applications poses strong demands on the HPC storage subsystem, since data needs to be typically copied from compute nodes to I/O nodes and vice versa when calculations start and stop. In this scenario, the emerging trend of adding denser, NVM-based storage to compute nodes as burst buffers, offers the possibility of using these resources to construct temporary filesystems that perform ad-hoc I/O optimizations for specific batch jobs. Thus, we introduce the concept of a temporary filesystem called echofs that coordinates with the job scheduler to preload a job's data dependencies into a collaborative burst buffer, which is created by virtually aggregating the NVM burst buffers available to compute nodes. The filesystem distributes data across nodes with a pseudo-randomized striping strategy to favor load balance, offer aggregated I/O and allow for efficient data redistributions when the number of compute nodes change between jobs. In addition, echofs offers a POSIX interface and internally manages the access to the different burst buffers available to a compute node (e.g. SSDs or NVM), so that researchers and legacy applications can interact with the filesystem without significant issues.

5.3 Exploring User Level Burst Buffer on Public Cloud and HPC

Tianqi Xu (Tokyo Institute of Technology, JP)

As the big data grows, more and more computational resources are required to run large scale applications, However, the strict access in HPC centers prevents public users from running their applications at large scale. In this talk, we introduce CloudBB, a user-level and on-demand filesystem to exploit burst buffer techniques for clouds. With CloudBB, large scale applications can be executed on public cloud instead of HPC for easier access, and virtually unlimited resource. Unlike conventional filesystems, CloudBB creates an on-demand two-level hierarchical storage system and caches popular files to accelerate I/O performance. CloudBB enables scalable I/O with multiple metadata servers, and also is resilient to failures by using file replication, failure detection and recovery techniques. On the other hand, in recent HPC systems, the computational performance has been rapidly increasing. However, the performance of parallel file systems in HPC has not been able to keep up with computing power. This limits the performance of scientific data-intensive applications in HPC. In addition, parallel file systems are shared by all the users, which exacerbates the situation due to I/O contention. We also introduce HuronFS, an extension of CloudBB to solve the problems in HPC. By building dedicated burst buffer on-demand, we can accelerate applications and avoid I/O contention. HuronFS exploits low-latency and high-bandwidth interconnects (InfiniBand) for HPC systems. Our evaluation shows HuronFS can achieve up to 3.5 GB/sec per a single I/O node in I/O throughput, which is comparable to the average throughput of the parallel file systems in Tokyo Tech's supercomputer, TSUBAME2.5.

5.4 MarFS and DeltaFS at LANL: User-Level File Systems Opportunities and Challenges

Brad Settlemyer (Los Alamos National Laboratory, US)

In this talk I discuss MarFS and DeltaFS, two user-level file systems efforts at Los Alamos National Laboratory that provide greater flexibility in how larg high performance computing (HPC) data centers procure storage systems. MarFS leverages a POSIX namespace and distributed object storage technology to construct a scalable long-term storage system that is capable of storing extremely large scientific data sets. The architecture enables a subset of the POSIX consistency model relevant to long term storage use cases while also supporting simple capacity addition over time. DeltaFS provides an extremely scalable metadata plane that decouples application metadata load from the underlying HPC platform storage system. Both systems allow LANL to procure storage systems with great flexibility while supporting ultrascale scientific simulation workloads.

5.5 Gfarm file system meets burst buffer

Osamu Tatebe (University of Tsukuba, JP)

File system for data processing like Google FS, HDFS and Gfarm, exploits local access performance and avoids network data transfer. This is especially relevant now with node local NVMe SSD and NVRAM. But there are some changes too, more performance required, durabilitly not as important in some modes, and data may be temporary. This talk designs Gfarm for burst buffers to implement RDMA access and no durability option.

5.6 HPC File System Opportunities with Microkernels (LWK), especially L4-based

Hermann Härtig (TU Dresden, DE)

I have tried to point out to the file system experts present what micro kernels can do for the construction of user-level file systems. In HPC-operating systems research, a variant of systems has become popular again after a period of hibernation: a micro kernel(like L4) and a monolithic OS (like Linux) are combined and applications are split in two parts one running on the micro kernel and the other on Linux. A number of interesting questions then arises such as: who is the boss (L4 or Linux), what is the programming model for the micro kernel, how to simplify the programming of the split application. I described how in a combination of L4 and Linux, some of the questions are answered: 1) in contrast to competitors L4 has a message-passing/memory-management oriented interface, much better suited for user-level memory-management than the common a Posix-like interface, 2) L4 is in control over Linux, which help isolate critical applications from irregularities of Linux, 3) Linux applications can simply leave the sphere of control of Linux, freely float on the available cores, and be pushed back if Linux SysCalls are needed. I left it the FS experts present to decide how to use that opportunities.

6 Talks Session: Object Stores and Alternatives

The fourth session of the Dagstuhl seminar focused on research and development of object stores as an alternative self-describing storage solution to popular but wieldy parallel file systems. This session features four separated talks that covered the employment of object stores for distinctly different use cases.

Suren Byna from Lawrence Berkeley National Laboratory in the U.S. started the session with his talk on **Proactive Data Containers (PDC) for next generation HPC storage**. This work is from an ASCR research project funded by the U.S. Department of Energy to explore next generation storage systems and interfaces. In view of the hefty impedance imposed by the deep storage hierarchy to vertical data movements, Suren and his colleagues introduced PDC as the new storage paradigm that offers an object-centric data abstraction for storage management on next-generation exascale systems. PDC is particularly keen on addressing the fundamental performance challenges caused by the consistency requirements and stateful semantics of the POSIX I/O calls. To this end, the PDC team has introduced a number of novel features as part of the PDC storage architecture, including a simple client API for light-weight object manipulation, SoMeta for scalable metadata management and a data elevator for automatic and transparent data movement across storage hierarchies. All these features are implemented as part of the HDF5 I/O model for scientific datasets.

Andreas Dilger from the Intel High Performance Data Division presented the second talk on **DAOS: A Scale-Out Object Store for NVRAM and NVMe**. The Distributed Asynchronous Object Storage (DAOS) project is sponsored by the U.S. Department of Energy. Different from the PDC effort, the DAOS project is targeted as a general storage stack to exploit low-latency byte-granular NVRAM and NVMe storage for any upper-level storage paradigms such as MPI-IO, ADIOS and even PDC. DAOS leverages proven technologies for its communication protocol and fine-grained lockless client I/O. Similarly to PDC, it also adopts the container abstraction as an integral part of its five-tier storage model for a novel object store. These tiers simplify the formation of automatic storage workflows for applications and enable transparent encapsulation of a wide variety of different storage interfaces. As Andreas indicated, the overarching goal of the project is to overcome POSIX limitations and offer a unified storage model over which domain-specific data models can be developed, such as HDF5, ADIOS, HDFS and Spark, in the meantime, increasing data velocity by several orders of magnitude. The project is in active community development. The current code has been released open-source on GitHub (https://github.com/daosstack/daos).

Maria S. Perez from the Universidad Politëcnica de Madrid and Luc Bougë from INRIA of France jointly gave the third talk on **Are objects the right level of abstraction to enable the convergence between HPC and Big Data at storage level?**. Given the dual demands from HPC and data analytics applications, Maria and her team have developed an object based storage system called Tyr that aims to provide a converged storage architecture. Maria indicated that Tyr is designed with scalable concurrency and synchronization management. Using BLOBs (Binary Large OBjects) as the basic building blocks, Tyr enables the store and retrieval of unstructured data through simple object-based binary methods. Tyr internally maps all writes to transactions. Its transactions are not limited to a chunk for a file, but may span multiple chunks and blocks. The WARP algorithm is employed for high-performance, sequentially-consistent transaction management. Maria and her team have evaluated Tyr using the MonALISA data collection and comparing with other object-based storage systems including Ceph, Azure and BlobSeer. Her results showed that Tyr delivers good performance advantages for transactional writes and reads. Furthermore, Maria believes that Tyr is well suited for both HPC and Big Data Analytics applications, for which a number of nice figures with comparative results are also included in her talk.

Finally, Lance Evans and Raghunath Raja Chandrasekar from Cray Inc jointly gave a talk on **An Exploration into Object Storage**. In this talk, they presented their exploration of a converged storage architecture for both HPC simulation and analytics platforms. Cray's position is that a converged storage platform shall emerge through hardware and software co-design of exascale-class supercomputers and analytics frameworks. A solution they are currently exploring is a library called SAROJA (Scalable And Resilient ObJect storAge). In its software architecture, SAROJA consists of three main components. The first is its user-space API library designed a host of intelligent mechanisms including algorithmic selection of data nodes, automatic translation of posix commands to KV/NoSQL operations. The other two components are the meatadata service and data path services. Lance and Raja designed these two critical pieces in a decoupled manner to allow flexible evolution. Particularly, Lance and Raja explored the use of Casandra NoSQL operations for the metadata services, and a number of Ceph components for data path services. They claim to have seen promising results from the prototype, which warrants future development and plenty of opportunity for advancement.

6.1 Proactive Data Containers (PDC) for next generation HPC storage

Suren Byna (Lawrence Berkeley National Laboratory, US)

Emerging high performance computing (HPC) systems are expected to be deployed with an unprecedented level of complexity, due to a multi-layer memory/storage hierarchy. This hierarchy is expected to range from CPU cache through several levels of volatile memory to non-volatile memory, traditional hard disks, and tape. Simple and efficient methods of data management and movement through this hierarchy is critical for scientific applications using exascale systems. Existing storage system and I/O technologies face severe challenges in dealing with these requirements. POSIX and MPI I/O standards that are the basis for existing I/O libraries and parallel file systems present fundamental challenges in the areas of scalable metadata operations, semantics-based data movement performance tuning, asynchronous operation, and support for scalable consistency of distributed operations.

Moving toward new paradigms for next-generation storage in the extreme-scale era, we have started investigating novel object-centric data abstractions and storage mechanisms that take advantage of the deep storage hierarchy and enable proactive automated performance tuning. We are developing a fundamentally new data abstraction, called Proactive Data Containers (PDC). A PDC is a container within a level of storage (memory, NVRAM, disk, etc.) that stores science data in an object-centric manner. In this talk, I will present an evolution of parallel I/O paradigms, high-level libraries such as HDF5, and our recent progress in PDC system development, PDC application programming interface (API), scalable metadata management and data movement optimizations aimed for object-centric storage.

6.2 DAOS: A Scale-Out Object Store for NVRAM and NVMe

Andreas Dilger

The Distributed Asynchronous Object Storage (DAOS) is an open-source (https://github.com/daos-stack/daos) storage stack designed from the ground up to exploit low-latency byte-granular NVRAM and NVMe storage. DAOS provides true byte-granular data and metadata I/O using persistent memory, combined with NVMe storage for bulk data. This, combined with end-to-end OS bypass, results in ultra-low I/O latency. The DAOS stack aims at increasing data velocity by several orders of magnitude over conventional storage stacks, while providing extreme scalability and resilience.

The DAOS API abstracts the underlying multi-tier storage architecture and offers a unified storage model over which domain-specific data models can be developed, such as HDF5, ADIOS, HDFS and Spark. Legacy POSIX access is similarly built on top of the DAOS

API. The core data model is a byte-granular key-value store which allows I/O middleware to overcome POSIX limitations and to have access to advanced capabilities like non-blocking I/O, ad-hoc concurrency control, distributed snapshots, native producer-consumer pipeline, end-to-end data integrity, index/query and in-situ analysis. DAOS also provides scalable, distributed transactions to I/O middleware with improved data consistency and automated recovery.

6.3 Are objects the right level of abstraction to enable the convergence between HPC and Big Data at storage level?

Maria S. Perez (Universidad Politëcnica de Madrid) and Luc Bougë (INRIA – Rennes, FR)

One of the most challenging aspects of storage in HPC is addressing concurrent write sharing. This scenario appears in applications requiring ACID transaction support or having situations in which concurrent updates may lead to incoherent results. Many HPC storage systems rely on locking to coordinate the access when concurrent writes are needed. This locking mechanism introduces a non negligible performance overhead.

On the other hand, storage systems are converging towards the use of a distributed object storage model, as a way of providing scalability and more flexibility. BLOBs (Binary Large OBjects) appear in this arena as a very good alternative. They allow storing unstructured data accessed through low-level binary methods.

However, the scalability and high performance of BLOBS are achieved at the expense of weak consistency. Tyr is a new BLOB storage architecture, which provides efficient multichunk and multiblob transactions under heavy access concurrency, but guaranteeing sequential consistency. A significant advantage of this approach is that it enables the building of lightweight file systems on top of this layer. Moreover, due to its features, Tyr can be used as underlying storage system for both HPC and Big Data Analytics applications, contributing to the desired convergence between these two worlds.

6.4 An Exploration into Object Storage

Lance Evans (Cray Inc. – Seattle, US) and Raghunath Raja Chandrasekar (Cray Inc. – Seattle, US)

The need for scalable, resilient, high performance storage is greater now than ever, in high performance computing. This talk presents exploratory research at Cray that studies aspects of emerging storage hardware and software design for exascale-class supercomputers, analytics frameworks, and commodity clusters. Our outlook toward object storage and scalable database technologies is improving as trends, opportunities, and challenges of transitioning to them also evolve. Cray's prototype SAROJA (Scalable And Resilient ObJect storAge) library is presented as one example of our exploration, highlighting design principles guided by the I/O semantics of HPC codes and the characteristics of up-and-coming storage media. SAROJA is extensible I/O middleware that has been designed ground-up with object semantics exposed via APIs to applications, while supporting a variety of pluggable file and object back-ends. It decouples the metadata and data paths, allowing for independent implementation, management, and scaling of each. Initial functional and performance evaluations indicate there is both promise and plenty of opportunity for advancement.

7 Talks Session: File Systems Building Blocks

The development process of a (parallel) file system typically takes ten years, before a file system matures. User-space file systems should, according to the previous sessions, be able to address specific application demands, which requires a short design cycle.

Shorter times to mature can partly be achieved by moving the file system functionality from kernel space to user space and therefore enable better debugging capabilities as well as access to libraries with predefined functionalities. One aspect of this session is therefore the necessary library support, which can be used in many user-level file system implementations. A second aspect to reduce the development time is to make file systems *simpler* by removing functionality, which is not required for *all* applications. The simplifications can include waiving POSIX guarantees, but also to generally relax security via capabilities or ACLs (if it can be guaranteed by system architecture).

The talk **Security Issues in User-Level File Systems** from Stergios Anastasiadis focused on the current state of the art on how to provide security guarantees within our user-space file system scenario. He started with a general overview on datacenter multitenancy, where virtual machines access disk images rather than files in an Infrastructure as a Service. The traditional ways of federated access control in IaaS include centralized, Peer-to-peer, or Mapping-based (e.g., HekaFS) approaches. Sharing of co-located data resources is typically inflexible due to the lack of appropriate sharing architecture (with Amazon Elastic File System being a notable exception, while still relying on NFS that limits scalability and security).

The Dike-approach, developed in Stergio's group, provides mechanisms for hierarchical identification and authentication, where the provider manages tenants and the tenants manage users. The native multi-tenant authorization separate ACLs per tenant and provider and therefore allows for an efficient permission management, including shared common permissions and inheritance. The Dike prototype has been implemented over Ceph.

Security for user-space files systems has to include, according to Stergios, standard aspects like confidentiality, integrity, and availability, flexibility concerning the policies, efficiency, as well as support for local and remote devices as well as for byte-addressable *and* block-based access. Efficient access to SCM includes that the policy mechanisms cannot rely on kernel mediation, as the switch to the kernel space already removes the latency advantages of SCM. Proposed approaches in literature to include security in untrusted environments include Secure containers (SCONE, OSDI 2016), limited kernel mediation via Arrakis (OSDI 2014), file systems for SCM (Aerie, EuroSys 2014), and remote direct storage access (DiDAFS, HotOS 2015).

SCONE secure containers are based on Intel's Software Guard Extension SGX that provides secure regions of user-mode address spaces as enclaves, which are associated with an enclave page cache, where data from the enclave page cache is encrypted and integrity-protected when sent to memory. SCONE provides secure containers by protecting the confidentiality and integrity of application data even against an adversary with super-user access to the operating system and hardware. All internal data is protected through SGX, while SCONE provides transparent encryption and authentication of data via shields, e.g., to send system calls to the host OS through shared memory or for writing to the file system. SCONE distinguishes between unprotected, authenticated, encrypted and authenticated files, while an ephemeral file system maintains state in non-enclave memory.

The main idea of Arrakis is to split the operating system tasks in two different entities. The kernel operates in the control plane and has to configure the hardware to limit application misbehavior. Applications in Arrakis can then independently and directly access virtual I/O devices and allows most I/O operations to be performed without prior kernel access. Arrakis implements a hardware-independent layer for virtualized I/O, including a virtual interface controller, doorbells, filters, and rate specifier.

Aerie proposes a flexible file system, where user-mode programs can directly access storage-class memory without kernel interaction. Aerie is based on a decentralized architecture. lifFS is an untrusted library providing a file-system interface including naming and data access via persistent memory primitives, which is linked to the user space application. A centralized trusted file-system service enforces metadata integrity and synchronization by providing distributed lock service and the SCM Manager is a kernel component for SCM allocation, mapping and protection through hardware privileges by providing memory permissions and virtual address mapping.

DiDAFS provides remote allocation mechanisms of main memory or storage device areas to support remote direct memory (and storage) access. The control plane of the device-owning node therefore exposes possible allocations to the network cards, which then can initiate RDMA accesses to remote user virtual memory. A user-level file system is used to manage metadata, while the metadata negotiation can still be handled through standard RPC mechanisms. Data, which is not cache-resident at the data accessing node, can therefore be mapped and efficiently fetched.

Stergios's overview has shown that security for user-space file systems still includes open issues when coupling the proposed ideas concerning protection (limited mediation, virtualization), consistency (processor cache, local/remote caches), interfaces (POSIX, memory-mapped, key-value, high-level data structures), and communication (library, RPC, RDMA, software fault isolation).

The motivation for Federico Padua's talk **An effort to systematically understand UFS design requirements** has been to understand the important design points when starting the development of a user-space file system and which file system components need to be implemented and which components can be thrown away. Furthermore he wants to understand whether tools can help to characterize application and link their characteristics to suggestions for FS component design.

Key file system components are data and metadata management, file system caches, fault tolerance, concurrency control and consistency as well as permissions and related security mechanisms. Especially concurrency control, consistency, and security seem (according to Federico) to be barely touched in file system research as long as a file system is not intended to be practically used.

Federico's talk is then mostly focused on file locking and further concurrency control mechanisms, as today's parallel file systems Lustre and GPFS (as POSIX files systems) implement quite complex distributed byte-range locking protocols to safely access data. Lustre and GPFS employ a pessimistic approach to concurrency control where a conflict is always expected, so that the file systems lock all accesses. OrangeFS on the contrary uses an optimistic approach, where a conflict is not expected and is only dealt with when it (rarely) occurs. POSIX.1-2008 on the other side does not specify the behavior of concurrent writes to a file from multiple processes, as applications should use some form of internal concurrency control.

The POSIX approach poses the question, whether file locking is that important in HPC at the file system level and whether concurrent accesses should better be correctly handled in applications or libraries (and whether they do it). Serialization (and therefore locking) might become necessary in case of overlapping writes to regions of a file or even to the same byte, in cases of clients reading after data has been written before, so that the client reads the last update. Furthermore, locking can happen in case of false sharing, where two processes from different nodes try to write to different bytes in a file mapping to the same block, or in case of writing to internal file system data structures.

Federico then presented his current work to measure file locking activity to determine the concurrency control mechanism needed in a file system. Starting from a pool of 12 applications from bioinformatics, chemistry, soft matter physics, high-energy physics, and weather modeling he traces in a first step the locking activity (using GPFS tracing) before characterizing applications to understand whether locking has been necessary.

Michael Kuhn's talk **The Case for a Flexible HPC Storage Framework** is motivated by the observation that it is hard to investigate new file system approaches, as file systems are typically monolithic in design and researchers have to change many different components of a file system to investigate a new idea. This leads, according to Michael, to two major problems. First, many specialized solutions do exist for a particular problem, which are seldom contributed back and second, it is necessary to have a complete understanding of a file system before a researcher can contribute to file system research, which is an unnecessary hurdle for young researchers and students.

An additional trend is that many applications rely on high-level I/O-libraries like NetCDF or HDF5, which support the exchange of data based on their self-describing properties. Tightly coupling these interfaces with current file systems is hard to achieve and is most often being performed in the context of big projects like DAOS or ESiWACE. Similar trends can be seen in the context of HPC and big data convergence or of alternative file system interfaces.

The main focus of Michael's talk has been on the design of JULEA, which is a framework to support rapid prototyping of new ideas via plugins for interfaces, storage backends, or new semantics. JULEA should help to overcome the burden that many projects have to re-implement basic functionality from scratch before being able to focus in core innovations. JULEA runs completely in user space and supports plugins that are configurable at runtime. High-level libraries and applications can use it directly, so that it provides a convenient framework for research and teaching.

JULEA therefore makes it possible to offer arbitrary interfaces to applications, both traditional file system interfaces and completely new ones. The possibility to change client implementations or the storage backend should foster experimentation, while it is even possible to dynamically adapt the semantics, e.g., between POSIX and MPI-IO on a per-operation basis. JULEA therefore does not put any restrictions on the client interface and its user-space implementation removes many restrictions of the standard kernel-space VFS architecture. This is, according to Michael, useful for applications and I/O libraries, as libraries like HDF5 can be placed directly on top of JULEA.

The server backends are separated into data and metadata backends. Additionally, client and server backends. Data backends manage objects and their design is influenced by file systems (Lustre and OrangeFS), object stores (Ceph's RADOS), and I/O interfaces (MPI-IO).

Metadata backends manage key-value pairs and are influenced by databases (SQLite and MongoDB) and key-value solutions (LevelDB and LMDB).

An aim of JULEA is it to make file system adapt to applications instead of the other way around. The semantics concerning atomicity, concurrency, consistency, ordering, persistency, and safety of operations can be changed at runtime and it is possible to mix the settings for each of these categories, even if not all combinations might produce reasonable results.

JULEA has been implemented in C11 as an open source project, which can be accessed at https://github.com/wr-hamburg/julea. It integrated support for tracing and unit tests. Future improvements will include an HDF5 VOL plugin, a data backend for Ceph's RADOS and a metadata backend for LMDB.

Phil Carns talk **Why should we still talk about data service building blocks?** is also discussing ways to simplify file system design, but takes a slightly different approach than Michael's recommendations. He started with a reflection about the origins of conventional storage service architectures, which have been based on the key technologies of block devices, sockets, pthreads, and kernel drivers. There have only been few cores per node and the concurrency has been further reduced by introducing I/O forwarding nodes. Data has been served via dedicated, remote service nodes, which have been either connected through TCP/IP sockets or in HPC over dedicated high-speed networks. The costs to tie the different components have not mattered too much, as the underlying magnetic disks had service times in the order of milliseconds.

An updated version of the same storage service architecture shows some key changes. New technologies like NVRAM, RDMA, dynamic services, and a much higher concurrency have been putting much higher pressure on keeping latencies and jitter low. Technology transitions include the shift from block devices to byte-addressable memory, moving from sockets to small messages and RDMA as well as the introduction of dynamic service groups and cluster-local services. Phil's main message is to encourage researchers to not implement everything again and again but to share engineering components and to free up time for research.

Phil therefore presented an example data path which is based on open source components like Argobots, Margo, Mercury, CGI, and Libpmem. The resulting architectural diagram looks like including many layers, but Phil argues that the number of components is not a major challenge as long as the interactions between the layers are efficient. The keys to optimize interaction between layers is to avoid privileged mode transitions, context switches in general, and memory copies.

The formerly mentioned Libpmem is a user space library written by Intel to access persistent memory regions (NVML). It provides well-defined control over persistency and device naming/reference conventions. A family of derived libraries for data structures that understand persistent memory references, transactions, or atomicity have been developed, like Libpmemobj for object storage, Libpmemblk accessing fixed-size blocks, Libvmmalloc, which replaces the standard `malloc()` command, or Pmemfile as a file system in user space with no kernel VFS mediation.

Mercury is an RPC system for use in the development of high performance system services, which has been developed by the HDF Group and ANL. It is portable across systems and network technologies and provides efficient bulk data movement to complement control messages. It has been built on lessons learned from, e.g., IOFSL, Nessie, or lnet (see https://mercury-hpc.github.io). It sits on top of transport libraries using a plugin API for network abstractions and supports OFI, IB, TCP/IP, and SHMEM transport protocols. Mercury provides simplifications for service implementers like remote procedure calls, RDMA abstraction, protocol encoding, or clearly defined progress and event models without requiring a global fault domain like `MPI_COMM_WORLD`. Phil mentioned that Mercury is minimal and fast, but a challenge to program. Therefore his group has added a layer We added a layer called *Margo* that adds an easy-to-use sequential interface, which relies on user-level thread scheduling for concurrency.

An earlier observation in Phil's talk has been that there is a much higher I/O concurrency than there are cores available on storage service nodes. Each of the requests can be illustrated as a state machine, where each request involves multiple steps: some will block, some will branch and the server needs to keep track of where each request is in state machine. ANL has therefore introduced ARGOBOTS as a user-level threading framework with lightweight context switching among many concurrent threads. Key features for data services are that it enables developers to track the state of many concurrent operations with simple service code paths and low OS resource consumption. ARGOBOTS also allows custom schedulers to implement priorities or limit CPU usage and provides primitives that facilitate linkage to external resources.

Phil closed his talk with an example of how to glue multiple components together when providing a huge set of microservices. Composed libraries can, e.g., exchange RPCs with a composed service, where the composed service delegates RPCs to microservices and propagates RDMA tokens to the service(s) that will drive data transfer. Generally, this approach allows to mix and match remote and non-remote components with the same API conventions, while not changing the service implementation. In a last step, services can be combined in scalable service groups (SSGs) (see https://xgitlab.cels.anl.gov/sds/ssg), which add group membership to Mercury by bootstrapping groups, providing identifiers to concisely reference groups of processes, and providing optional fault detection.

The philosophical question whether **To FS or not to FS...** has been asked by Rob Ross. He experienced that user level approaches are typically employed as a light weight and flexible way to provide functionality that does not have a viable business model driving vendor support. Functionality is therefore limited to just what is needed to accomplish the task at hand. However, the potential for such user level services enables new, fundamental services that might become keystones of future HPC systems. Main questions are how do the particular I/O use cases inform the way we manage data, how we should present hierarchical storage systems to user applications and how should we manage data movement through a storage hierarchy.

The specialization of data services can already be seen analyzing executables and libraries like SPINDLE, checkpoint-restart tools like SCR and FTI, as well as intermediate data products like DataSpaces, Kelpie, or MDHIM, which all are not filesystems. Many of them are provisioned via MPI or the batch system, use local storage, but only provide minimum fault management and security.

Rob argues in favor of Phil's approach when demanding for an ecosystem of data services, where many components can be shared across multiple services. Some of these services will look like file systems, while others take a completely different form. To build such an ecosystem, it is important to tackle hard problems like group membership, authentication and authorization, publish/subscribe **and** performance. These basic services will, according to Rob, enable a broader community to build better, more capable user-level data services than possible today.

The idea behind such an ecosystem is presented based on the example of future applications exploring the use of multi-scale modeling. The application Lulesh, e.g., loosely couples continuum scale models with more realistic constitutive/response properties. Fine scale model results can be cached and new values interpolated from similar prior model calculations. Goals of a simulation run are to minimize costly fine scale model executions, query/response times to databases and to load balance accesses to the distributed database. A possible approach is to start with a key/value store, to distribute approx. nearest-neighbor queries, to distribute data to co-locate values for interpolation and to import/export results to persistent stores. All these steps do not have to include file systems.

A closing remark from Rob included the observation that many parallel file systems today are veneers on object stores. While users consistently complain about file systems that they cannot create as many files (names) as they want and are not happy with metadata management, these are not the problem of the underlying object stores. Some users have therefore moved to structured APIs like HDF5, ADIOS, or netCDF and many classes of data are used where the file system is not visible to the user at all. Interestingly, object stores themselves are not controversial and probably the best bet for the organization of long(er) term storage.

7.1 Security Issues in User-Level File Systems

Stergios V. Anastasiadis (University of Ioannina, GR)

We provide a brief overview of our recent research activities on reducing the interference from compactions in sorted key-value data stores, improving the performance of object-based distributed filesystems with host-side journaling, and adding native multitenancy support to object-based distributed filesystems. Then we summarize recent design points on secure user-level filesystems and identify possible open issues for further research on isolation, caching, interface and communication.

7.2 An effort to systematically understand UFS design requirements

Federico Padua (Universität Mainz, DE)

File locking is a form of serialization used by parallel filesystems to prevent data corruption or other bugs in case many programs or many processes access the same files at the same time. This is the case for a single MPI program, in which many processes read and write a single shared file or multiple files during the run. If user-level filesystems aim to replace some key functionalities of kernel level filesystems, then it becomes important to verify whether such a filesystem needs to support any form of concurrency control or none. In case a form of concurrency control is needed, it's also important to asses which one would be the best option. Our interest in studying concurrent file accesses and file locking is motivated by one section of the POSIX specification: "This volume of POSIX.1-2008 does not specify behavior of concurrent writes to a file from multiple processes. Applications should use some form of concurrency control.". It then becomes fundamental that filesystem designers take into account this possible issue when developing their filesystems: either some form of concurrency control is needed or not. This mandates a clear understanding and study of real world applications that run in HPC centers.

We hope that the study we are embarking on will open up a discussion and a rethinking of some assumptions. For instance, previous works tried to explore the benefits of optimistic concurrency control in the HPC context. We plan first to assess to what extent ad-hoc file system developers should care about concurrency control or not, to support one design choice over another: this can justify optimistic or pessimistic approaches.

In the case of more general purpose filesystems in HPC clusters, such as GPFS and Lustre, the default to concurrency control is to employ a pessimistic approach using some kind of distributed locking (with different level of granularity). We plan to use the study to provide a solid experimental foundation to see whether this pessimistic approach is needed even in the case of general purpose filesystems to correctly support HPC applications.

7.3 The Case for a Flexible HPC Storage Framework

Michael Kuhn (Universität Hamburg, DE)

JULEA is a flexible storage framework that allows the offering of arbitrary client interfaces to applications. To be able to rapidly prototype new approaches, it offers data and metadata backends that can either be client-side or server-side; backends for popular storage technologies such as POSIX, LevelDB and MongoDB have already been implemented. Additionally, JULEA allows dynamically adapting the I/O operations' semantics and can thus be adjusted to different use-cases. It runs completely in user space, which eases development and debugging. Its goal is to provide a solid foundation for storage research and teaching.

7.4 Building blocks for user-level HPC storage services

Philip Carns (Argonne National Laboratory, US)

Today's most well-known distributed storage systems were envisioned and developed in an era dominated by Linux clusters and magnetic disks, and the technology building blocks that went into those storage systems were well understood: TCP/IP, local files or block devices, POSIX threads, and OS kernel drivers. New HPC system trends are putting pressure on limitations of those building blocks, however. Storage device latency is dropping, CPU clock speed advances are slowing, and emerging applications demand flexibility in interfaces and provisioning. If our building blocks don't evolve to reflect these trends, then system software overheads will soon become the number one limiting factor in storage system efficiency.

In this talk I will discuss how technology building block assumptions have changed in recent years and the challenges and opportunities that have arisen as a result. In particular I would like to highlight how RDMA-enabled network layers, persistent memory interfaces, lightweight threads, and user-level interfaces are taking the place of TCP/IP, files and blocks, POSIX threads, and kernel modules in HPC storage system design.

7.5 To FS or not to FS...

Robert B. Ross (Argonne National Laboratory, US)

In the HPC community, user level approaches are typically employed as a light weight and flexible way to provide functionality that caters to specific applications or that doesn't have a viable business model driving vendor support. In this scenario, functionality is limited to just what is needed to accomplish the task at hand. However, the potential for user level services goes far beyond the examples that we see today, potentially providing new, fundamental services that are keystones of future HPC systems. In this talk I will discuss some aspects of distributed systems that HPC user level service developers often avoid, highlighting some areas for collaboration on (or leveraging of) components and micro services that would open up new opportunities for user level services in HPC environments.

8 Discussions

8.1 Discussion: How are stage in and stage out operations actually going to work? *Lead: Ned Bass*

On upcoming systems, we assume we will be able to move data to and from the tier of storage on or near the compute nodes before and after a job in stage-in and stage-out operations. We assume that the operations will improve performance, because jobs will not have to wait for the potentially expensive data transfer operations before they begin or after they end. Given all that, it is unclear how these operations will work in realistic situations. In this session we dove into questions surrounding the practical implementation and implications of stage-in and stage-out operations.

Are these stage-in and -out operations part of the job or do they belong to the resource manager?

- We will have to pay the cost somehow. How do we save what we need and minimize the impact of data movement?
- We need a way for users to define, specify what they want for particular data files. What is the information that the resource manager needs? What is the API?
- Are these operations performed during the run? Or pre-/post-run? What happens if pre-stage is supposed to happen before the run but the operation is not complete and the nodes are ready for the job? Should the job be launched anyway?
- What are the security concerns when someone else's data is on the storage but someone else is running?
- Are you consuming network bandwidth that could impact other jobs with stage-in/out? Is the impact on other jobs big enough to offset the wasted time due to waiting for files to transfer (idle compute nodes)

What does this mean from a scheduler perspective?

- It is complex from a resource scheduling perspective.
- Not only is it scheduling resources for data movement (network, capacity), but it is also trying to schedule nodes taking into account additional resources.

- Scheduling jobs is much harder with node-local storage case. Multi-constraint scheduling problem
- Trade-off between the time to wait for jobs to run while data is fetched and the time to read data from the parallel file system. Which is worse?

Are the questions about this different for shared burst buffers versus local burst buffers?

- Does stage-in/out affect jobs with local burst buffers more or less than with shared burst buffers?
- It's still a problem of shared network resources, storage resources
- Perturbation is unavoidable in either case
- Does this mean we need to schedule I/O? E.g. Delay checkpoint transfer to prestage data for the next job, then trickle down checkpoint? Is scheduling too big of a hammer here? What about prioritization instead?
- Could intermix jobs that need I/O with those that don't when possible.

How can a center figure out its priorities and policies for stage-in and out?

- A multi-objective optimization problem between utilization, job interference, performance
- Could we have a dedicated storage network? That would be awesome but expensive

8.2 Discussion: How can we fairly judge the performance of storage systems – IO 500? *Lead: John Bent*

This discussion centered around how to define a set of tests for evaluating storage systems fairly. A good number of participants had already been working towards this with the IO 500, so the discussion was in the context of the IO 500. Evaluating storage systems is more challenging now that we have hierarchical storage systems and a wider variety of important I/O workloads. The IO 500 working group has already been working to define a set of tests; this discussion session introduced the tests and solicited feedback from the seminar participants.

What kinds of tests should we use?

- Easy test: What is the best you can do with your system? Danger of this is that people will just submit their results for the fastest tier of storage, e.g. burst buffers.
- Hard tests – comparisons across systems. Which one has better characteristics for different workloads?
- What is the right way to collect the measurements for a single comparable metric?
- Do we account for performance changes as systems age?
- Bandwidth tests, metadata tests
- A variety of tests for different I/O workloads including well-behaved and poorly-behaved applications. Real applications have complex behavior. We need a good set of proxy applications to capture this.
- What about a real application mix for evaluation? That would be informative for purchasing. However, real applications are very cumbersome and can take a long time. Emulators would be sufficient if they stress the I/O system in the same way.

How do we future-proof the tests?

- We don't want the defined tests to be out of date with the next generation of systems
- Want to make it independent on node size, size of operations
- Should we use relative metrics then?

What metric(s) should we use to evaluate the systems?

- Is a single metric adequate? You really want the details, number of clients, etc in some cases. Could the metric be a list of values for different layers of storage? We will want to know the performance of the layers individually as well as in aggregate
- The metric definition that will be most valuable differs given the situation at hand. A conflict between vendors and procurers
- A single metric is appealing though due to human psychology. People will want to know how they are ranked. Once they look for that, they will look at the details
- Would be good if centers could contribute their acceptance testing criteria for storage systems. Perhaps we could derive a metric from those. They could tell us what is hard/important for their workloads.

8.3 Discussion: What is a user level file system? What do we mean when we say that? *Lead: Philip Carns*

In this discussion, we set out to define what is meant when we say 'user level file system' as it became evident that we all had different definitions in our minds.

What do we mean by user level? Is it just bypassing the OS? And what does that mean?

- Do we mean no system calls to do the I/O or networking?
- No data copies necessary between user and kernel space?
- Is there a difference between user-level and user-space file system? No.
- Can a user deploy the file system or does it take an administrator?

Should we be talking about storage systems instead of file systems?

- Probably yes so we can include object stores
- Well, if we're storing data persistently then perhaps it is sufficient to call it a file system. What do we mean by persistent? Nothing lasts forever, e.g. Lustre purges. That is due to center policies, not Lustre.
- Let's settle on ULSS – user level storage systems since it seems to include everything

Why do we want user level storage systems?

- Easily deployed by user, deployed quickly
- Does not require root/administrator
- Limited lifetime file system, quick to spin up/tear down. Does not consume resources when not needed.
- Ownership of file system is user
- Developed and run at the user level, without kernel hacking
- Failure isolation
- Tailored to specific workloads, customizable. Can use different consistency models. Every metric by which you evaluate storage can be tailored: performance, durability, etc
- OS bypass: is this a good thing? A requirement? Is the only reason we want OS bypass because privilege escalation and de-escalation is high latency?
- With OS bypass we can access the raw device

What are the trade-offs for performance with user-level versus traditional file systems?

- Can add transformative, transparent improvements in performance due to customizability
- Can essentially give the user a “consistency” button, tuner. They can determine their own durability, persistence needs
- Users can choose the implementation that is best tuned to their use case, possibly violating POSIX but it is fine for their needs

8.4 Discussion: How can we characterize what users need from storage systems? *Lead: Scott Klasky*

The key question for this discussion was how can we as a community understand what users need from storage systems? If we do not understand what they need, then our efforts towards optimization will be useless. We need to have a way to characterize the I/O behavior in enough detail to get the information we need.

Can we define a set of patterns or “motifs” that describe application I/O behavior?

- If we can decide what the motifs are, then can we generate benchmarks that are reasonably representative?
- How do we decide if the motifs/patterns are accurate?
- What if they change over time? For example, we have data analytics workloads emerging now.
- What characterization tools can we use to learn the motifs? A single application may contain multiple motifs. For example, checkpoint/restart and output
- What if we define motif differently so that instead of being an I/O type it is more of an application type? The question here is do all applications of the same “type” do the same kind of I/O?

Can we build up a cache of mini apps for I/O behaviors?

- How would we generate the mini apps? From data or from an emulator like IOR or MACSio?
- What do we do about too many mini apps? Different across organizations?
- How do we capture the differences in behavior with scaling, input, architecture?
- We need to track all the I/O types from an application, the behavior of which will be use case and system dependent. For example, the frequency of checkpointing will change depending on (perceived) failure rate
- How will we validate these mini apps?
- We need to be sure to state the purpose of the mini apps. A mini app can’t do everything. Need to clearly define what the mini app is intended to mimic.

8.5 Discussion: Are we ready to program to a memory hierarchy versus block devices? *Lead: Jay Lofstead*

This section documents the discussion led by Jay Lofstead surrounding the question of whether the storage and I/O community is ready for changes in hardware expected on upcoming systems, namely working with extended memory devices as opposed to block devices. In general, because this represents a potential paradigm shift for not only the storage and I/O community but also for the programming models community, this discussion brought up more questions than it did answers.

What is memory and what is storage?

- Memory that is slow may need to be treated by programmer as storage. However, people are already doing out of core computation and treating storage as slow memory.
- If it's temporary data, then it's memory. If it's meant to persist, then it's storage.
- Conclusion seems to be that the difference is in how you use the device more than the speed of the device.

What do applications people need?

- What is familiar to applications people? Is it dealing with memory?
- Upcoming machines have burst buffers, but we will have machines that only have devices that can be addressed as memory very soon.
- Legacy software could continue to do traditional I/O, but perhaps new applications or smaller applications will change to a new programming model. Need to support both models.
- Does this question belong to the programming models folks? Not sure.
- How do we help users figure out what to store and what to recompute?

How long do we store data at each tier?

- Depends on capacity, devices near compute do not have much capacity
- Energy cost of moving data, too high compared to compute. It's much cheaper to move data across compute nodes than to move it to the parallel file system and back
- Do we want to provide hierarchical storage management software/model that users can specify high level needs of data? Can users describe the high level needs for their data in a workflow (where to move it, how long to keep it)?
- Storage hierarchy levels are not independent – data can exist in several locations, with newer versions on compute node storage, older versions on parallel file system. The metadata becomes much more complicated than it is today.
- How can users find their data in such a complex system?

How do we characterize user workloads?

- There are very many types of user workloads from experimental data analysis, but at the 50k foot view they are very similar. HPC workloads tend to write some data during a time step.

  There are different workloads than that though, e.g. ensemble applications, data analysis applications. In general we can think of them as bulk output and asynchronous I/O patterns.
- What will be the requirements of data analytics workloads? Random I/O? No workload characterizations of these yet.

  Random I/O is hard to support – either it fits in the storage tier you are using or it doesn't.
- We really need to have good characterizations of the applications to understand what they are doing

Will we be able to get away from application specific frameworks that solve a specific problem?

- ADIOS, HDF5 provide application-specific interfaces that have proven very useful in some domains

- If we have higher level application interfaces that map to some common API (e.g., POSIX) is that the answer?
- The cloud world offers many interfaces/services for different workloads, seems to be a better idea than trying to shoehorn in an application into a particular I/O model
- In a layered model (application interface over lower-level common interface) they key will be to maintain interoperability
- Can we renegotiate the contract with users on what the interface is? Google never asks. Well, in HPC users are the boss.

Can we make changes to POSIX or define a POSIX-like interface with relaxed constraints?

- Do users understand POSIX? Do we even understand all of POSIX? No.
- Applications developers will not change their I/O approach unless we demonstrate clear, strong benefit. The benefit of POSIX is that it is stable, well benefit and hindrance.
- Relaxing the constraints of POSIX could be a powerful approach

8.6 Discussion: What should we do about POSIX? *Leads: Rob Ross and Andreas Dilger*

The group discussed what to do about the POSIX interface going forward. It has drawbacks and can prevent I/O performance due to its strict semantics. However, it is widely used and has been stable for a very long time.

Do we have a hope of getting away from POSIX?

- There are a mass of tools and applications that already use POSIX, legacy codes. Can we expect them to change?
- Are we stuck with POSIX as long as we have parallel file systems?

How successful are middleware libraries at getting around the problems of POSIX?

- The difficulties of implementing ROMIO on top of POSIX really exposed the problems of POSIX and that implementations of POSIX do not uphold the contract of the API
- Log structured optimizations preclude being able to read the file from POSIX
- Can we write object stores on top of POSIX? There are not products that do this now, but there could be. However it is a semantic mismatch.

Can we keep POSIX and relax the semantics in a defined way? A new contract?

- APIs are a contract, but can we define how to relax the semantics? Eg. relax the semantics on the burst buffer but uphold them on the parallel file system?
- Perhaps upholding POSIX semantics is easier in the burst buffer because of fast compute network
- POSIX is not good for coupled applications communicating through the file system
- Maybe POSIX could be the contract between the I/O middleware and the backend store. We could give applications something else to program to without the strict semantics of POSIX
- POSIX is going to be around a long time. We need a new way to express our storage requirements with or without POSIX
- Perhaps as a community we could come to a consensus on key ways to extend and relax POSIX, e.g. NFSv4 semantics

9 Recurring Themes and Future Work from Seminar

Overall, the week of the seminar was used productively. The group enjoyed the balance of short talks and discussion sessions. Since this was the first time that this group had come together, it was beneficial to have the short talks to give everyone an overview of people's backgrounds and expertise. The primary results of the seminar were identification of needs that the HPC I/O community should address in the approach to exascale computing. These needs arose as recurring themes throughout the seminar week.

9.1 Recurring Themes

Throughout the talks and in-depth discussions in the seminar, the participants identified several recurring themes that represent items that need to be addressed by the HPC storage and I/O community in the coming years as we come to terms with hierarchical storage systems. The needs that were identified during the seminar were:

- The need to characterize I/O behavior of applications
- The need to re-evaluate the POSIX I/O interface
- The need to advance resource management and scheduling policies for I/O management
- The need to define users' data management requirements.

The need to characterize the I/O behavior of applications. Seminar participants identified the need to characterize the I/O behavior of applications in several sessions. In general, the participants felt that as a community we don't have a good grasp on the aspects of application I/O behavior that are needed to design high-performance storage systems. For example, are HPC applications well-behaved in that if they write shared files, does each file offset only get updated by one process? Or do the applications assume that the file system implements locking and thus the writes of different processes to the same offset will have some ordering enforced? The answers to these questions will inform storage system design. In the case of this example, if applications are well-behaved, then a storage system will not need to implement expensive locking and thus I/O operations can have better performance. However, if the application depends on locks for coherent file data, then the storage systems will have to support that.

Quite a bit of discussion in the seminar centered on how to generate a set of mini-apps that mimic true I/O behavior. Much of the work in mini-app design is centered around other aspects of application behavior, e.g., computation or communication patterns, and hardly any attention is given to I/O behavior. Of the few I/O mini-apps that do exist, for the most part, the I/O patterns that are emulated are too abstracted to be of use for next-generation storage system design and simply focus on bulk write and read phases. We need to find or develop mini-apps that mimic different application I/O patterns and are true to varying behavior based on different inputs and scaling of runs. We need to ensure that we collect mini-apps that are representative of all major classes of application I/O behavior, from classic bulk-synchronous checkpoint/restart, to newer I/O models for in-situ analysis or machine-learning workloads.

As part of the I/O characterization effort, we also need to focus on measurement of real applications running on systems. This measurement effort can aid in the generation of useful I/O mini-apps, but also provide additional information for designing storage systems and scheduling jobs and I/O operations on future systems. While several quality I/O measurement tools exist today, I/O measurement is usually done as an ad hoc activity, measuring the I/O operations of a particular application to understand its behavior and performance. While this is helpful for that particular application on the system it was run on, it does not give a picture of how that application's performance was affected by concurrently running jobs that may have been using network or storage resources, and it does not capture workflow-level I/O interactions where a series of application runs over a long-running, months-long experiment may share file data. Thus, it will be useful to have detailed I/O measurement of applications on a center-wide basis and to couple that with storage system and resource management information. Several centers are making headway in this regard currently, but given that workloads tend to vary across HPC centers, it will be beneficial to have more centers invest in in gathering broad-spectrum I/O information from jobs.

The need to re-evaluate the POSIX I/O interface. The topic of POSIX came up quite often during the week of the seminar. Participant opinions spanned from full rejection and replacement of the ubiquitous standard to acceptance and possible redefinition of semantics thought to be too restrictive for today's and next-generation HPC storage systems. In general, most participants agreed that the POSIX I/O interface is too widely used and relied upon by legacy HPC applications to imagine its disappearance any time soon. Additionally, as a community, we do not have an alternative to POSIX that is as stable and portable to offer application developers.

In general, participants did not envision a near-future end to POSIX, and many felt that there was an opportunity to possibly alter the semantics of POSIX to be less restrictive in order to better support HPC workloads. This effort would require understanding the needs of HPC I/O workloads with respect to file integrity, which further emphasizes the need for comprehensive I/O characterization and measurement as discussed above. However, even before full understanding of application I/O behavior is known, as a community we have a sense of the limitations of POSIX semantics and can make an effort to systematically identify these limitations and define structured relaxations to the POSIX I/O semantics. For example, well-behaved applications that write shared files may be able to specify that locking is not needed by providing a defined flag to the interface. Participants of the seminar plan to begin a community effort for this need.

The need to advance resource management and scheduling policies for I/O management. With hierarchical storage systems becoming a reality for many centers, the consensus of the seminar participants was that resource manager and scheduling policies have not been updated to account for them. Current production resource managers and schedulers only account for the computational requirements of jobs, the node count and run time. However, in order to achieve better system utilization and throughput, resource managers and schedulers should account for new storage system models. For example, node-local storage devices like burst buffers may need to be allocated and shared across jobs (the currently executing job, the previous job that may be draining data, and the next job to run that may have data pre-staging into the node). All of these will require resources including storage capacity, possibly some CPU time depending on the storage device, and network resources. If resource managers and schedulers are able to allocate resources and schedule jobs to minimize the interference of competing I/O activities, then the result will be higher overall system performance.

There will be significant additional complexity in accounting for the resources needed for I/O in hierarchical storage systems. As a community, we plan to work with researchers in the area of resource management to provide input on how storage and network resources might be used on future systems in order to prepare.

The need to define users' data management requirements. User applications are reading and writing data at volumes greater than ever and the trend promises to continue in the future. This trend towards large data volumes was the impetus behind hierarchical storage systems, where storage devices closer to compute nodes could act as write-through caches or buffers for the I/O operations between stable storage (e.g., the parallel file system) and compute nodes, or alternatively to use the cache tier as temporary storage for data that only has a temporary life span and does not need to be transferred to the parallel file system. The canonical example of the temporary use case is checkpoint/restart, where checkpoint files can quickly be written to the cache tier only, and upon writing a new checkpoint file, an older checkpoint can be deleted since it will not be needed anymore. Another example of using the cache tier as a temporary storage layer is of an in-situ analysis component coupled with a simulation application. In this case, the simulation could output large data files on the cache tier, which could be analyzed in-situ by the analysis component. Following analysis, the files could either be deleted if not needed further, or perhaps only "interesting" output files (as determined by the analysis component) could be transferred to the parallel file system.

In general, hierarchical storage systems offer new options for I/O operations of HPC applications, but there is no general way for applications to specify how a particular data set should be managed on the storage tiers. Examples of information that could be specified include whether the data is temporary or if it should be persisted on stable storage; the requirements on resilience of the data, where perhaps the user can afford to lose some data files with a specified probability; and time before data must be persisted on stable storage, where a user might optionally allow a very slow drain of particular data if it is not needed immediately.

As a community, we agree that we should come together with use cases and requirements in order to develop a specification for users to provide the needs of their data sets. The specification will enable storage system developers to provide higher performance implementations because they will not need to implement the most rigorous case of persisting all output data on stable storage.

9.2 Future Work

The participants of the seminar tended to agree that our community should engage in more meetings of the style we experienced in this Dagstuhl Seminar. The discussion-oriented meeting was fruitful in teasing out overarching themes that represent future work items that need to be addressed by our community. The participants indicated that future seminars could be centered around designing community-supported solutions to one or more of the themes that we identified during this seminar. The group also indicated that they would like to participate in a future joint publication effort, e.g., a special issue journal containing papers from our collective efforts related to this seminar to serve as a basis for future research.

Participants

- Stergios V. Anastasiadis
University of Ioannina, GR
- Ned Bass
LLNL – Livermore, US
- John Bent
Seagate Government Solutions – Herndon, US
- Thomas Bönisch
HLRS – Stuttgart, DE
- Luc Bougë
INRIA – Rennes, FR
- Andrë Brinkmann
Universität Mainz, DE
- Suren Byna
Lawrence Berkeley National Laboratory, US
- Philip Carns
Argonne National Laboratory, US
- Lance Evans
Cray Inc. – Seattle, US
- Wolfgang Frings
Jülich Supercomputing Centre, DE
- Hermann Härtig
TU Dresden, DE
- Dean Hildebrand
IBM Almaden Center – San Jose, US
- Nathan Hjelm
Los Alamos National Laboratory, US
- Scott Klasky
Oak Ridge National Laboratory, US
- Michael Kluge
TU Dresden, DE
- Michael Kuhn
Universität Hamburg, DE
- Jay Lofstead
Sandia National Labs – Albuquerque, US
- Satoshi Matsuoka
Tokyo Institute of Technology, JP
- Alberto Miranda
Barcelona Supercomputing Center, ES
- Kathryn Mohror
LLNL – Livermore, US
- David Montoya
Los Alamos National Lab., US
- Federico Padua
Universität Mainz, DE
- Mark Parsons
University of Edinburgh, GB
- Maria S. Perez
Universidad Politëcnica de Madrid, ES
- Franz Josef Pfreundt
Fraunhofer ITWM – Kaiserslautern, DE
- Raghunath Raja Chandrasekar
Cray Inc. – Seattle, US
- Robert B. Ross
Argonne National Laboratory, US
- Brad Settlemyer
Los Alamos National Laboratory, US
- Osamu Tatebe
University of Tsukuba, JP
- Tianqi Xu
Tokyo Institute of Technology, JP
- Weikuan Yu
Florida State University – Tallahassee, US
- Yue Zhu
Florida State University – Tallahassee, US

Report from Dagstuhl Seminar 17221

Geometric Modelling, Interoperability and New Challenges

Edited by

Falai Chen[1], Tor Dokken[2], Thomas A. Grandine[3], and Géraldine Morin[4]

1 **Univ. of Science & Technology of China – Anhui, CN,** chenfl@ustc.edu.cn
2 **SINTEF – Oslo, NO,** tor.dokken@sintef.no
3 **The Boeing Company – Seattle, US,** thomas.a.grandine@boeing.com
4 **University of Toulouse, FR,** geraldine.morin@enseeiht.fr

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17221 "Geometric Modelling, Interoperability and New Challenges". While previous Dagstuhl seminars on geometric modeling were focused on basic research, this seminar was focused on applications of geometric modeling to four topic areas: big data and cloud computing, multi-material additive manufacturing, isogeometric analysis, and design optimization. For this purpose we brought together participants from industry urgently in need of better solutions, researchers in the above application areas, and researchers in the geometric modeling community.

Seminar May 28–2, 2017 – http://www.dagstuhl.de/17221
1998 ACM Subject Classification G.1 Numerical Analysis, G.4 Mathematical Software, I.2 Artificial Intelligence, I.3 Computer Graphics, J.2 Physical Sciences and Engineering
Keywords and phrases additive manufacturing, computer graphics, design optimization, geometric modeling, geometry, geometry processing, isogeometric analysis, shape design
Digital Object Identifier 10.4230/DagRep.7.5.140
Edited in cooperation with Georg Muntingh

1 Executive Summary

Falai Chen
Tor Dokken
Thomas A. Grandine
Géraldine Morin

This report documents the program and the outcomes of Dagstuhl Seminar 17221 "Geometric Modelling, Interoperability and New Challenges".

The importance of accurate geometric models of shapes, both naturally occurring and man-made, is rapidly growing with the maturing of novel manufacturing technologies and novel analysis technologies. The advent of big data challenges and cloud-based computing serves to confound the distribution and remote access of these geometric models. While previous Dagstuhl seminars on geometric modeling were focused on basic research, this seminar was focused on applications of geometric modeling. We selected four core application areas that stretch the underlying mathematical underpinnings of the discipline to its limits and beyond:

- Big data and Cloud computing
- Multi-material additive manufacturing (3D Printing)
- Isogeometric analysis
- Design optimization

The seminar provided a forum for leading researchers to present new ideas, to exchange scientific insights, and to bring together practical applications and basic research. The previous seminar explored the theory of shape representations, shape transformations, and computational models for each. This seminar explored application of this theory to the four above-mentioned application challenges, strengthening the reliability and performance of applications in engineering, manufacturing, and scientific exploration. The goal of the seminar was to establish a common understanding between the Geometric Modeling research community and the above application fields by addressing the following questions:

- How to handle geometric descriptions that cannot be processed by one computer but have to be distributed among many computers in the cloud?
- How to process huge 3D data sets that may be noisy and incomplete into clean, scalable, and easy to access 3D environments?
- How to turn big, dispersed maybe noisy and incomplete geometric data into clean, scalable, and easy to access 3D information that can be used for change detection and decision making?
- How to represent, control, and process complex, anisotropic, internal material enabled by additive manufacturing, and how to design for additive manufacturing?
- How to perform topology optimization of the internal structure of objects to enable additive manufacturing to reach its full potential?
- How to introduce analysis-based design in CAD systems with isogeometric analysis in mind?

To answer these questions we brought together participants from industry urgently in need of better solutions, researchers in the application areas represented by the four topics, and researchers in the geometric modeling community whose interests align with the four topic areas. The scientific presentations lasted 15 to 20 minutes. Senior researchers gave 4 overview talks on the 4 themes of the seminar. Perspective working groups were organized for each of the four topics.

2 Table of Contents

Executive Summary
Falai Chen, Tor Dokken, Thomas A. Grandine, and Géraldine Morin 140

Overview of Talks

Planar Pythagorean-Hodograph B–spline Curves
Gudrun Albrecht 145

The Role of Volume Meshes in 3D Modeling for Additive Manufacturing
Marco Attene 145

Big Data Impact on Geometric Modeling
Chandrajit Bajaj 146

Low Rank Approximation for Planar Domain Parameterization
Falai Chen 146

Introductory Talk on Additive Manufacturing
Tor Dokken 147

Precise Construction of Micro-structures and Porous Geometry via Functional Composition
Gershon Elber 147

Isogeometric Finite Element Methods for Shape Optimization
Daniela Fußeder 148

THB-spline Simplification and Scattered Data Approximation
Carlotta Giannelli 148

Visualization of Uncertainty-aware Geometry
Christina Gillmann 149

De Boor Suitable T-splines
Ron Goldman 149

Making Multidisciplinary Design Optimization Work
Thomas A. Grandine 149

Drall-based Ruled Surface Modeling
Hans Hagen and Benjamin Karer 150

Shape from Sensors
Stefanie Hahmann 150

An Error Bound for Interpolation by Low Rank Functions
Bert Jüttler 151

Shape Preserving Interpolation on Surfaces
Panagiotis Kaklis and Alexandros Ginnis 151

Truncated Hierarchical Loop Subdivision Surfaces and Application in Isogeometric Analysis
Hongmei Kang, Falai Chen, and Xin Li 152

Branched Covering Surfaces
Konrad Polthier (Freie Universität Berlin) 152

Assembly Design of Wind-up Toys
Ligang Liu 152

Generalized Splines on T-meshes
Tom Lyche 153

Spectral Analysis of Matrices from NURBS Isogeometric Methods
Carla Manni 153

Multi-scale Differential Analysis of Complex Shapes
Nicolas Mellado 154

Navigation in 3D Networked Virtual Environments: Sprite Trees and 3D Bookmarks
Géraldine Morin 155

B-spline-like Bases for C^2 Cubics on the Powell-Sabin 12-Split
Georg Muntingh and Tom Lyche 155

T-junctions in Spline Surfaces
Jörg Peters 156

Isogeometric Analysis on Triangulations
Xiaoping Qian 156

Moments of Sets with Refinable Boundary
Ulrich Reif 156

Weights for Stabilised Least Squares Fitting
Malcolm A. Sabin 157

Spline-based Quadrature Schemes in Isogeometric Analysis for Boundary Element Methods
Maria Lucia Sampoli, Francesco Calabrò, and Giancarlo Sangalli 157

Introductory Talk on Isogeometric Analysis
Giancarlo Sangalli 158

Interactive Deformation of Virtual Paper
Camille Schreck 158

A Framework for Isogeometric Analysis on Unstructured Quadrilateral Meshes
Hendrik Speleers 159

Using Machine Learning Techniques in Geometric Modeling
Georg Umlauf 159

A Multi-sided Bézier Patch with a Simple Control Structure
Tamas Várady 160

Shape Optimization for Fine-scale Structure Design for 3D Printing
Denis Zorin 160

Working groups

Additive Manufacturing Focus Group
Tor Dokken, Thomas A. Grandine, and Georg Muntingh 161

Isogeometric Analysis Focus Group
Tor Dokken, Thomas A. Grandine, and Georg Muntingh 162

Big Data Focus Group
Thomas A. Grandine, Tor Dokken, and Georg Muntingh 164

Open problems

Technical position paper on trivariates
Malcolm A. Sabin 166

Participants 168

3 Overview of Talks

3.1 Planar Pythagorean-Hodograph B–spline Curves

Gudrun Albrecht (Universidad Nacional de Colombia – Medellin, CO)

Joint work of Gudrun Albrecht, Carolina Vittoria Beccari, Jean-Charles Canonne, Lucia Romani

We introduce the new class of planar Pythagorean-Hodograph (PH) B–Spline curves. They can be seen as a generalization of the well-known class of planar Pythagorean-Hodograph (PH) Bézier curves, presented by R. Farouki and T. Sakkalis in 1990, including the latter ones as special cases. Pythagorean-Hodograph B–Spline curves are nonuniform parametric B–Spline curves whose arc-length is a B–Spline function as well. An important consequence of this special property is that the offsets of Pythagorean-Hodograph B–Spline curves are non-uniform rational B–Spline (NURBS) curves. Thus, although Pythagorean-Hodograph B–Spline curves have fewer degrees of freedom than general B–Spline curves of the same degree, they offer unique advantages for computer-aided design and manufacturing, robotics, motion control, path planning, computer graphics, animation, and related fields. After providing a general definition for this new class of planar parametric curves, we present useful formulae for their construction and discuss their remarkable attractive properties. Then we provide a method to determine within the set of all PH B–Splines the one that is closest to a given reference spline having the same degree and knot partition.

3.2 The Role of Volume Meshes in 3D Modeling for Additive Manufacturing

Marco Attene (CNR – Genova, IT)

Joint work of Daniela Cabiddu, Marco Attene
Main reference D. Cabiddu, M. Attene, "ϵ-maps: Characterizing, detecting and thickening thin features in geometric models", Computers & Graphics, Vol. 66, pp. 143–153, Elsevier, 2017.
URL https://doi.org/10.1016/j.cag.2017.05.014

Classical product development pipelines include a design phase and a process planning that preceeds the actual fabrication. When moving from design to process planning, smooth NURBs-based surfaces are typically tessellated, and the resulting triangle meshes are analyzed and prepared for the final slicing step. This talk highlights the limitations of this surface-based approach, and discusses how explicit volume meshes can be exploited to effectively resolve a number of tasks, including the geometric analysis to ensure compatibility with the terget printer, and the proper description of objects constituted of multiple and graded materials.

References

1 D. Cabiddu, and M. Attene, *Epsilon-maps: Characterizing, detecting and thickening thin features in geometric models.* Computers&Graphics (Elsevier, in press).

3.3 Big Data Impact on Geometric Modeling

Chandrajit Bajaj (University of Texas – Austin, US)

Joint work of Benedikt Bauer, Radhakrishna Bettadapura, Antje Vollrath, Abhishek Bhowmick, Eshan Chattopadhyay, David Zuckerman, Andrew Gillette, Alexander Rand
URL https://cvcweb.ices.utexas.edu/cvcwp/publications/

We revisit classic problems of geometric modeling in this age of "Big Data" and present scalable solutions with accuracy / speed tradeoffs. The scalable techniques include dimension reduction of spline spaces, generating low-discrepancy samplings of product spaces with quasi-polynomial number of samples (as opposed to exponential in the product dimension), approximate non-uniform fast Fourier transforms, shape optimized dictionaries, and approximate geometric optimization. The latter is by reduction to Semi-Definite Programming (SDP), and methods for generating almost regular, congruent tiled arrangements using a new generative class of polyhedra. Applications of these scalable techniques are to functional data and shape approximation, geometric shape similarity, complementarity matching and constructing multi-component assemblies.

References

1 Chandrajit L. Bajaj, Benedikt Bauer, Radhakrishna Bettadapura, Antje Vollrath: *Nonuniform Fourier Transforms for Rigid-Body and Multidimensional Rotational Correlations.* SIAM J. Scientific Computing 35(4) (2013)
2 Chandrajit L. Bajaj, Abhishek Bhowmick, Eshan Chattopadhyay, David Zuckerman, *On Low Discrepancy Samplings in Product Spaces of Motion Groups.* CoRR abs/1411.7753 (2014)
3 Andrew Gillette, Alexander Rand, Chandrajit Bajaj, *Construction of Scalar and Vector Finite Element Families on Polygonal and Polyhedral Meshes.* Comput. Meth. in Appl. Math. 16(4):667–683 (2016)

3.4 Low Rank Approximation for Planar Domain Parameterization

Falai Chen (Univ. of Science & Technology of China – Anhui, CN)

Joint work of Maodong Pan, Falai Chen

Construction of spline surfaces from given boundary curves is one of the classical problems in computer aided geometric design, which regains much attention in iso-geometric analysis in recent years and is called domain parameterization. However, for most of the state-of-the-art parametrization methods, the rank of the spline parameterization is usually large for complex computational domains, which results in higher computational cost in solving numerical PDEs. In this talk, we propose a low-rank representation for the spline parameterization of planar domains using low-rank tensor approximation technique, and apply quasiconformal mapping as the framework of the spline parameterization. Under given correspondence of boundary curves, a quasi-conformal map with low rank and low distortion between a unit square and the computational domain can be obtained by solving a non-linear optimization problem. We propose an efficient algorithm to compute the quasi-conformal map by solving two convex optimization problems alternatively. Experimental results show that our approach can produce a low-rank parametric spline representation of planar domains.

References

1 Bert Juettler and Dominik Mokriš, *Low rank interpolation of boundary spline curves*. Computer Aided Geometric Design, 2017, vol.55, 48–68.

2 Xianshun Nian and Falai Chen, *Planar domain parameterization for isogeometric analysis based on teichmüller mapping*. Computer Methods in Applied Mechanics and Engineering 2016, vol.311, 41–55.

3.5 Introductory Talk on Additive Manufacturing

Tor Dokken (SINTEF – Oslo, NO)

URL http://www.caxman.eu

The first part of the talk addressed Additive Manufacturing (AM) from the perspective of manufacturing. Differences between AM, subtractive manufacturing and formative manufacturing were explained. The evolution of AM was outlined, fundamentals of AM explained including the formal definition of AM by ISO/ASTM 52900. A short description of the seven fundamentally different AM process categories provided. In the second part of the talk the digital perspective of AM was in focus exemplified by the EC H2020 R&I Action CAxMan – Computer Aided Technologies for Additive Manufacturingm. In CAxMan a 3-variate isogeometric approach for simulation based design targeting AM is addressed. CAxMan has a strong focus on the challenges of interoperability of design, simulation, process planning and actual manufacturing for AM, and development of suitable standards supporting AM.

3.6 Precise Construction of Micro-structures and Porous Geometry via Functional Composition

Gershon Elber (Technion – Haifa, IL)

We introduce a modeling constructor for micro-structures and porous geometry via curve-trivariate, surface-trivariate and trivariate-trivariate function (symbolic) compositions. By using 1-, 2- and 3-manifold based tiles and paving them multiple times inside the domain of a 3-manifold deforming trivariate function, smooth, precise and watertight, yet general, porous/micro-structure geometry might be constructed, via composition. The tiles are demonstrated to be either polygonal meshes, (a set of) Bézier or B-spline curves, (a set of) Bézier or B-spline (trimmed) surfaces, (a set of) Bézier or B-spline (trimmed) trivariates or any combination thereof, whereas the 3-manifold deforming function is either a Bézier or a B-spline trivariate. We briefly lay down the theoretical foundations, only to demonstrate the power of this modeling constructor in practice, and also present a few 3D printed tangible examples. We will then discuss these results and conclude with some future directions and limitations.

3.7 Isogeometric Finite Element Methods for Shape Optimization

Daniela Fußeder (MTU Aero Engines – München, DE)

Joint work of Bernd Simeon, Daniela Fußeder

We present a framework for shape optimization with isogeometric analysis (IGA) from an optimal control perspective: In this formulation both the state and the control space are discretized by B-splines which means practically that two (possibly different) meshes may be used, one for simulation and one for optimization.

As a result, interpolation errors of the state and the control may be addressed separately, yet through the connection of IGA a tight relation between geometry and analysis is maintained.

Moreover, several options like optimizing weights in a non-uniform rational B-spline geometry, using local refinement or exploiting higher polynomial degrees of basis functions are shown.

3.8 THB-spline Simplification and Scattered Data Approximation

Carlotta Giannelli (University of Firenze, IT)

Joint work of Cesare Bracco, Carlotta Giannelli, Alessandra Sestini

The geometric processing of unstructured large data sets may be effectively addressed by developing reliable adaptive schemes that guarantee high-quality approximations in terms of compact representations. By focusing on spline representations, adaptivity may easily be achieved by considering multilevel extensions of the standard B-spline model, where the tensor-product structure is (locally) preserved level by level.

We present two adaptive schemes that exploit the capabilities of truncated hierarchical B-splines (THB-splines) to reduce the computational costs connected with the reconstruction of large data sets. In particular, we introduce a THB-spline simplification algorithm defined in terms of an efficient data reduction operator. Subsequently, in order to deal with scattered data of high complexity and different nature, we present an adaptive data fitting scheme as extension of local least-squares approximations to hierarchical spline constructions. The perfomances of the two methods will be illustrated by a selection of examples, including real data sets describing different terrain configurations.

References

1 C. Bracco, C. Giannelli, and A. Sestini, *Adaptive scattered data fitting by extension of local approximations to hierarchical splines*. Computer Aided Geometric Design, 52-53:90–105, 2017.

3.9 Visualization of Uncertainty-aware Geometry

Christina Gillmann (TU Kaiserslautern, DE)

Joint work of Christina Gillmann, Thomas Wischgoll, Jose Tiberio Hernandez, Hans Hagen, Bernd Hamann

Images usually contain a variety of image errors such as lense flair, artifacts and blur. This errors are aggravating the application of image processing methods and reduce the quality of geometry extraction.To solve this problem, we present a novel description of geometry. In contrast to the classical definition of geometry we not solely consider points that can be composed to higher order objects. In addition each point obtains a support function descriping the possibility that a point exists and also a range where this point is allowed to be located. To be able to generate this geometry we also extend well known geometry extraction techniques to output the requested points and also their support functions. These support functions are based on uncertainty image metrics and can also be used to adapt the extracted geometry thus the resulting points obtain qualitatively higher support functions.

3.10 De Boor Suitable T-splines

Ron Goldman (Rice University – Houston, US)

Joint work of Yang Zhang, Visit Pataranutaporn, Ron Goldman

We describe necessary and sufficient conditions on the T-mesh for an Analysis Suitable T-spline to have a local de Boor recursive evaluation algorithm.

References

1. C. de Boor, *A practical guide to splines*, Volume 27 of Applied Mathematical Sciences, Springer-Verlag, New York, revised edition, 2001.
2. X. Li, *Some properties of analysis suitable T-splines*, Journal of Computational Mathematics, Vol. 33, No. 4, pp. 428–442.
3. X. Li, J. Zheng, T.W. Sederberg, T.J.R. Hughes, and M.A. Scott, *On linear independence of T-spline blending functions*. Computer Aided Geometric Design, Vol. 29, 2012, pp. 63–76.
4. Beirao da Veiga, A. Buffa, G. Sangalli, R. Vazquez, *Analysis-suitable Tsplines of arbitrary degree: definition and properties*, Mathematical Models and Methods in Applied Sciences, World Scientific Publishing Company.

3.11 Making Multidisciplinary Design Optimization Work

Thomas A. Grandine (The Boeing Company – Seattle, US)

In many engineering intensive design activities (e.g. airplanes and high performance automobiles), the shapes of the engineered products can have a tremendous effect on their engineering performance. Determining these shapes is often best accomplished by formulating optimization problems which attempt to maximize engineering performance subject to certain constraints. For example, one might wish to minimize the weight of an airplane subject to constraints on being able to carry a certain number of passengers a prescribed distance in a certain amount of time. Solving design problems this way imposes severe restrictions on geometric modeling methods as well as the analysis codes which determine that engineering performance. This talk will explore the interplay between optimization and the requirements it imposes on analysis and geometric design.

3.12 Drall-based Ruled Surface Modeling

Hans Hagen (TU Kaiserslautern, DE) and Benjamin Karer (TU Kaiserslautern, DE)

In continuum mechanics, ruled surfaces have shown considerable potential as candidates to modeling thin narrow strips of inextensible elastic material. Yet, the surface definitions usually applied in this domain are often limited in their generality or require complicated constraints to guarantee important geometrical properties. We introduce a model for the analytically exact definition of arbitrary ruled surfaces featuring trivial constraints for properties like the developability of a surface or the isometry of a deformation. The model is based on the movement of two coupled moving frames of reference whose transition equations are governed by a one-parametric minimal system of invariants of ruled surfaces. Based on these invariants, we derive a bending energy integral for arbitrary ruled surfaces of finite width. We demonstrate our model's descriptive power by several examples. Being general, accurate, and parallelizable, our model is well suited for high-detail modeling of arbitrarily shaped inextensible elastic surface strips.

3.13 Shape from Sensors

Stefanie Hahmann (INRIA Grenoble Rhône-Alpes, FR)

Joint work of Tibor Stanko, Stefanie Hahmann, Georges-Pierre Bonneau, Nathalie Saguin-Sprynski
Main reference T. Stanko, S. Hahmann, G.-P. Bonneau, N. Saguin-Sprynski, "Shape from sensors: Curve networks on surfaces from 3D orientations", Computers & Graphics, 66:74–84, Elsevier, 2017.
URL https://doi.org/10.1016/j.cag.2017.05.013

We present a novel framework for acquisition and reconstruction of 3D curves using orientations provided by inertial sensors. While the idea of sensor shape reconstruction is not new, we present the first method for creating well-connected networks with cell complex topology using only orientation and distance measurements and a set of user-defined constraints. By working directly with orientations, our method robustly resolves problems arising from data inconsistency and sensor noise. Although originally designed for reconstruction of physical shapes, the framework can be used for "sketching" new shapes directly in 3D space. We test the performance of the method using two types of acquisition devices: a standard smartphone, and a custom-made prototype for measuring orientations and distances.

References

1 Tibor Stanko and Stefanie Hahmann and Georges-Pierre Bonneau and Nathalie Saguin-Sprynski: *Shape from sensors: Curve networks on surfaces from 3D orientations*. (Shape Modeling International 2017), Computers and Graphics 66, pp.74–84, 2017.

3.14 An Error Bound for Interpolation by Low Rank Functions

Bert Jüttler (Universität Linz, AT)

It has been observed recently that tensor-product spline surfaces with low rank coefficients provide advantages for efficient numerical integration, which is important in the context of matrix assembly in isogeometric analysis [1]. By exploiting the low-rank structure one may efficiently perform multivariate integration by a executing a sequence of univariate quadrature operations. This fact has motivated us to study the problem of creating such surfaces from given boundary curves. More precisely, we propose a method for interpolation by rank-2 functions and apply it to the problem of creating surface patches from four given boundary curves [2]. In addition, we discuss the generalization to rank-n functions and analyze the relationship to the method of cross approximation [3]. In particular we derive a new error bound for interpolation by rank-n functions. Joint work with Nira Dyn and Dominik Mokriš.

References

1 A. Mantzaflaris, B. Jüttler, B.N. Khoromskij, U. Langer: Matrix Generation in Isogeometric Analysis by Low Rank Tensor Approximation. In: Curves and Surfaces, J.-D. Boissonat (ed.) et al., Springer LNCS, 2015, pp. 321–340.

2 D. Mokriš and B. Jüttler: Low rank interpolation of boundary spline curves. To appear in Comp. Aided Geom. Des. (2017).

3 J. Schneider: Error estimates for two-dimensional cross approximation. J. Approx. Theory 162 (2010), pp. 1685–1700.

3.15 Shape Preserving Interpolation on Surfaces

Panagiotis Kaklis (The University of Strathclyde – Glasgow, GB) and Alexandros Ginnis

We introduce a criterion for G2 shape-preserving interpolation of points on surfaces, based on the behaviour of the corresponding composite geodesic interpolant in the neighbourhood of the data points. Two alternative families of splines are proposed, combining appropriately nu- or non-uniform-degree splines with preimages of the geodesic segments, for conforming to the shape-preservation criterion for sufficiently large values of the nu or interval-degree parameters. Preliminary numerical results are given using the nu-splines-based interpolants for shape-preserving interpolation on cylinders, spheres and free-form surfaces.

3.16 Truncated Hierarchical Loop Subdivision Surfaces and Application in Isogeometric Analysis

Hongmei Kang (Univ. of Science & Technology of China – Anhui, CN), Falai Chen (Univ. of Science & Technology of China – Anhui, CN), and Xin Li

Joint work of Falai Chen, Hongmei Kang
Main reference H. Kang, X. Li, F. Chen, J. Deng, "Truncated Hierarchical Loop Subdivision Surfaces and application in isogeometric analysis", Computers & Mathematics with Applications, 72(8):2041–2055, Elsevier, 2016.
URL https://doi.org/10.1016/j.camwa.2016.06.045

Subdivision Surface provides an efficientway to represent free-formsurfaceswith arbitrary topology. Loop subdivision is a subdivision scheme for triangular meshes, which is C2 continuous except at a finite number of extraordinary vertices with G1 continuous. In this paper we propose the Truncated Hierarchical Loop Subdivision Surface (THLSS), which generalizes truncated hierarchical B-splines to arbitrary topological triangular meshes. THLSS basis functions are linearly independent, form a partition of unity, and are locally refinable. THLSS also preserves the geometry during adaptive h-refinement and thus inherits the surface continuity of Loop subdivision surface. Adaptive isogeometric analysis is performedwith the THLSS basis functions on several complexmodelswith extraordinary vertices to show the potential application of THLSS.

3.17 Branched Covering Surfaces

Konrad Polthier (Freie Universität Berlin)

Joint work of Konrad Polthier, Konstantin Poelke, Matthias Nieser, Ulrich Reitebuch, Faniry Razafindrazaka

Multivalued functions and differential forms naturally lead to the concept of covering surfaces and more generally of covering manifolds. This talk will review, illustrate and discretize basic concepts of branched covering surfaces starting from complex analysis, surface theory up to their recent appearance in geometry processing algorithms. Applications will touch artistic surface modeling, geometry retargeting, surface and volume parameterization, and novel weaved surface representations.

3.18 Assembly Design of Wind-up Toys

Ligang Liu (Univ. of Science & Technology of China – Anhui, CN)

Joint work of Ligang Liu, Peng Song, Xiaofei Wang, Xiao Tang, Chi-Wing Fu, Hongfei Xu , Niloy J. Mitra

Wind-up toys are mechanical assemblies that perform intriguing motions driven by a simple spring motor. Due to the limited motor forceand small body size, wind-up toys often employ higher pair joints of less frictional contacts and connector parts of nontrivial shapes to transfer motions. These unique characteristics make them hard to design and fabricate as compared to other automata. In this paper, we present a computational system to aid the design of wind-up toys, focusing on constructing a compact internal wind-up mechanism to realize user-requested part motions. Our key contributions include an analytical modeling of a wide variety of elemental mechanisms found in common wind-up toys, including their geometry and kinematics, conceptual design of wind-up mechanisms by computing motion transfer trees that support the requested part motions, automatic construction of wind-up mechanisms by connecting multiple elemental mechanisms, and an optimization on the geometry of parts and joints with an objective of compacting the mechanism, reducing its weight, and avoiding collision. We use our system to design wind-up toys of various forms, fabricate a number of them using 3D printing, and show the functionality of various results.

References

1 Peng Song, Xiaofei Wang, Xiao Tang, Chi-Wing Fu, Hongfei Xu, Ligang Liu, Niloy J. Mitra. *Computational Design of Wind-up Toys*. ACM Transaction on Graphics, 2017.

3.19 Generalized Splines on T-meshes

Tom Lyche (University of Oslo, NO)

Joint work of Bracco Cesaro, Tom Lyche, Carla Manni, Fabio Roman, Hendrik Speleers

Splines play a central role in IgA. We consider a generalization of polynomial splines known as Tchebycheffian splines. With such splines one can obtain nonrational representation of conic sections and Tchebycheffian B-splines have the important properties of the polynomial ones. Multivariate extensions of Tchebycheffian splines can be obtained via the tensor-product approach or by considering more general T-mesh structures. In this talk we consider Tchebycheffian spline spaces over planar T-meshes and analyze their dimensions, extending the results presented in [1, 2, 3].

References

1 C. Bracco, T. Lyche, C. Manni, F. Roman, and H. Speleers. *Generalized spline spaces over T- meshes: Dimension formula and locally refined generalized B-splines*. Appl. Math. Comput. 272 (2016), 187–198.
2 C. Bracco, T. Lyche, C. Manni, F. Roman, and H. Speleers. *On the dimension of Tchebycheffian spline spaces over planar T-meshes*. Comput. Aided Geom. Design, 45(2016), 151–173.
3 B. Mourrain. *On the dimension of spline spaces on planar T-meshes*. Math. Comp. 83 (2014), 847–871.

3.20 Spectral Analysis of Matrices from NURBS Isogeometric Methods

Carla Manni (University of Rome "Tor Vergata", IT)

Joint work of Carla Manni, Carlo Garoni, Stefano Serra-Capizzano, Hendrik Speleers

When discretizing a linear PDE by a linear numerical method, the computation of the numerical solution reduces to solving a linear system. The size of this system grows when we refine the discretization mesh. We are then in the presence of a sequence of linear systems with increasing size. It is usually observed in practice that the corresponding sequence of discretization matrices enjoys an asymptotic spectral distribution. Roughly speaking, this means that there exists a function, say f, such that the eigenvalues of the considered sequence of matrices behave like a sampling of f over an equispaced grid on the domain of f.

In this talk we analyze the spectral properties of discretization matrices arising from isogeometric Galerkin and collocation methods, based on d-variate NURBS of given degrees and applied to general second-order elliptic partial differential equations defined on a d-dimensional domain. This extends the results obtained for B-spline based isogeometric methods [1, 3].

The provided spectral information can be exploited for designing algorithms with convergence speed independent of the fineness parameters and also substantially independent of the degrees of the used NURBS, as in [2] for the B-spline case.

References

1 M. Donatelli, C. Garoni, C. Manni, S. Serra-Capizzano, and H. Speleers, *Spectral analysis and spectral symbol of matrices in isogeometric collocation methods*, Math. Comp. **85** (2016), 1639–1680.

2 M. Donatelli, C. Garoni, C. Manni, S. Serra-Capizzano, and H. Speleers, *Symbol-based multigrid methods for Galerkin B-spline isogeometric analysis*, SIAM J. Numer. Anal. **55** (2017), 31–62.

3 C. Garoni, C. Manni, S. Serra-Capizzano, D. Sesana, and H. Speleers, *Spectral analysis and spectral symbol of matrices in isogeometric Galerkin methods*, Math. Comp. **86** (2017), 1343–1373.

3.21 Multi-scale Differential Analysis of Complex Shapes

Nicolas Mellado (Paul Sabatier University – Toulouse, FR)

Joint work of Thibault Lejemble, Nicolas Mellado
Main reference N. Mellado, G. Guennebaud, P. Barla, P. Reuter, C. Schlick, "Growing Least Squares for the Analysis of Manifolds in Scale-Space", Computer Graphics Forum, 31(5):1691–1701, Wiley, 2012.
URL http://dx.doi.org/10.1111/j.1467-8659.2012.03174.x

With the proliferation of acquisition devices, gathering massive volumes of 3D data is now easy. Processing such large masses of point-clouds, however, remains a challenge, because of noise, varying sampling, missing data, and shapes complexity.

This talk will present a multi-scale shape analysis framework that aims at detecting arbitrary geometric features as local stability of the surface derivatives. Derivatives are obtained by differentiating implicit surfaces reconstructed at multiple scales from the point-cloud. Inspired by scale-space analysis, this framework is an attempt to provide similar analysis tools for unstructured 3d point-clouds, e.g. feature and pertinent scale extraction. This talk will also review how this framework can be used to improve point-cloud registration, shape abstraction and primitive approximation (point-cloud to CAD).

3.22 Navigation in 3D Networked Virtual Environments: Sprite Trees and 3D Bookmarks

Géraldine Morin (University of Toulouse, FR)

Joint work of Géraldine Morin, Minhui Zhu, Thomas Forgione, Wei Tsang Ooi, Axel Carlier, Vincent Charvillat
Main reference M. Zhu, G. Morin, V. Charvillat, W. T. Ooi, "Sprite tree: an efficient image-based representation for networked virtual environments", The Visual Computer, 33(11):1385–1402, Springer, 2017.
URL http://dx.doi.org/10.1007/s00371-016-1286-0

In this talk, we present how to adapt the 3D representation of large Networked Virtual Environments (NVE) to usage and navigation. Accessing and navigating through such large dataset is not easy for light clients with limited capacities, that is, with a limited rendering capacity, and a low bandwidth. First, we propose an alternative image based representation for the 3D content, generated by peers and collected in a sprite tree. We show that this representation leads to reducing both the data size but also the rendering time. Second, we show that we can further ease user navigation by proposing 3D bookmarks. These bookmarks also improve the predictability of the navigation: we benefit from the predictability in the streaming strategy and therefore can offer a better quality of service to the user.

References

1 Minhui Zhu, Géraldine Morin, Vincent Charvillat, Wei Tsang Ooi. *Sprite tree: an efficient image-based representation for networked virtual environments. Dans.* The Visual Computer, Springer-Verlag, Vol. (), p. 1–18, 2016. (first online)

3.23 B-spline-like Bases for C^2 Cubics on the Powell-Sabin 12-Split

Georg Muntingh (SINTEF – Oslo, NO) and Tom Lyche (University of Oslo, NO)

For spaces of constant, linear, and quadratic splines of maximal smoothness on the Powell-Sabin 12-split of a triangle, Cohen, Lyche and Riesenfeld recently discovered so-called S-bases. These are simplex spline bases with B-spline-like properties on a single macrotriangle, which are tied together across macrotriangles in a Bézier-like manner.

In this talk we give a formal definition of an S-basis in terms of certain basic properties. We proceed to investigate the existence of S-bases for the aforementioned spaces and additionally the cubic case, resulting in an exhaustive list. From their nature as simplex splines, we derive simple differentiation and recurrence formulas to other S-bases. We establish a Marsden identity that gives rise to various quasi-interpolants and domain points forming an intuitive control net, in terms of which conditions for C^0, C^1, and C^2-smoothness are derived.

Although the cubic bases can only be used to define smooth surfaces over specific triangulations, we envision applications for local constructions, such as hybrid meshes and extra-ordinary points, with the potential to be used in isogeometric analysis.

References

1 Tom Lyche and Georg Muntingh. *B-spline-like bases for C2 cubics on the Powell-Sabin 12-split.* Preprint.

3.24 T-junctions in Spline Surfaces

Jörg Peters (University of Florida – Gainesville, US)

Joint work of Kestutis Karciauskas, Daniele Panozzo, Jörg Peters

T-junctions occur in quad-meshes where surface strips start or terminate. While hierarchical splines are naturally suited for introducing T-junctions into quad meshes, they are not naturally suited for generating surfaces from quad meshes with T-junctions: some (many?) quad meshes with T-junctions do not admit a choice of knot-intervals for smooth hierarchical splines.

A simple geometric continuous spline construction, GT-splines, solves the problem, covering T-junctions by two or four patches of degree bi-4. GT-splines complement multi-sided surface constructions in generating free-form surfaces with adaptive layout.

References

1 Kestutis Karciauskas, Daniele Panozzo, Jörg Peters. *Spline surfaces with T-junctions.* https://arxiv.org/abs/1610.05351

3.25 Isogeometric Analysis on Triangulations

Xiaoping Qian (University of Wisconsin – Madison, US)

Joint work of Xiaoping Qian, Songtao Xia

In this talk, I will present a triangulation based isogeometric analysis approach. In this approach, both the geometry and physical fields are represented by bivariate/trivariate splines in Bernstein–Bézier form over the triangulation. We describe procedures for domain parameterization and triangulation from given physical domains, and procedures for constructing Cr-smooth basis functions over the domain. As a result, this approach can achieve automated meshing of objects of complex topologies, allow highly localized refinement, and obtain optimal convergence rates.

3.26 Moments of Sets with Refinable Boundary

Ulrich Reif (TU Darmstadt, DE)

Joint work of Jan Hakenberg, Ulrich Reif
Main reference J. Hakenberg, U. Reif, "On the volume of sets bounded by refinable functions", Applied Mathematics and Computation, 272:2–19, Elsevier, 2016.
URL https://doi.org/10.1016/j.amc.2015.06.050

We present a method for determining moments (volume, center of mass, inertia tensor, etc.) for objects which are bounded by refinable functions, such as 2d sets bounded by subdivision curves and 3d sets bounded by subdivision surfaces. The approach is based on the solution of an eigenequation resulting from the interrelation between a certain multilinear form and its refined counterparts.

3.27 Weights for Stabilised Least Squares Fitting

Malcolm A. Sabin (Numerical Geometry Ltd. – Cambridge, GB)

There is a set of well established ideas which seem to fit together in a way which is simple, powerful and general. Powerful because they seem to offer the combination of achieving desirable goals with a surprising level of possible automation. General because they seem to apply to a very wide subset of approximation theory applications. However, I have not yet seen these ideas presented in this combination before. Despite the generality available, the talk describes an experiment in the simplest possible context, that of scalar interpolation of univariate data, with promising results.

3.28 Spline-based Quadrature Schemes in Isogeometric Analysis for Boundary Element Methods

Maria Lucia Sampoli (University of Siena, IT), Francesco Calabrò, and Giancarlo Sangalli (University of Pavia, IT)

Joint work of Alessandra Aimi, Francesco Calabrò, Mauro Diligenti, Maria Lucia Sampoli, Giancarlo Sangalli, Alessandra Sestini

Boundary Element Methods (BEMs) are schemes studied since the mid '80, for the numerical solution of those Boundary Valued Problems, which can be transformed to Boundary Integral Equations. Indeed, if the fundamental solution of the differential operator is known, a wide class of elliptic, parabolic, hyperbolic, interior and exterior problems can be reformulated by integral equations defined on the boundary of the given domain, whose solution is successively obtained by collocation or Galerkin procedures. On the other side, the new Isogeometric analysis approach (IgA), establishes a strict relation between the geometry of the problem domain and the approximate solution representation, giving surprising computational advantages. In the IgA setting a new formulation of BEMs has been studied, where the discretization spaces are splines spaces represented in B-spline form ([1, 2]). In order to take all the possible benefits from using B-splines instead of Lagrangian basis, an important point is the development of specific new quadrature formulas for efficiently implementing the assembly phase of the method ([3]).

In this talk the problem of constructing appropriate and accurate quadrature rules, tailored on B-splines, for Boundary Integral Equations is addressed.

References

1 A. Aimi, M. Diligenti, M. L. Sampoli, A. Sestini; Isogeometric Analysis and Symmetric Galerkin BEM: a 2D numerical study, *Appl. Math. Comp.*, **272**, 173–186, (2016)
2 A. Aimi, M. Diligenti, M. L. Sampoli, A. Sestini; Non-polynomial spline alternatives in Isogeometric Symmetric Galerkin BEM, *Appl. Numer. Math*, **116**,
3 A. Aimi, F. Calabrò, M. Diligenti, M. L. Sampoli, G. Sangalli, A. Sestini, Efficient assembly based on B-spline tailored quadrature rules for the IgA-SGBEM, Submitted to *Computer Methods in Applied Mechanics and Engineering*, (under revision)

3.29 Introductory Talk on Isogeometric Analysis

Giancarlo Sangalli (University of Pavia, IT)

Joint work of P. Antolín, A. Buffa, A. Collin, G. Elber, M. Kapl, M. Martinelli, F. Massarwi, R. Vázquez, T. Takacs, M. Tani, et al.

This presentation is an overview of the main recent results on the numerical analysis and mathematical foundations of isogeometric analysis, introduced the seminal 2005 paper by Thomas J.R. Hughes, Austin Cottrell and Jury Bazilevs as a new approach to the discretization of partial differential equations (PDEs). Isogeometric analysis is a collection of methods that use splines, or some of their extensions such as NURBS (non-uniform rational B-splines), T-splines, LR-splines, hierarchical splines, etc. as functions to build approximation spaces which are then used to solve partial differential equations numerically.

The main challenge isogeometric analysis wants to address was to improve the interoperability between CAD and PDE solvers, and for this reason it uses CAD mathematical primitives, that is, splines and extensions, to represent PDE unknowns as well. Full interoperability is a challenging aim, for a number of reasons. For example, CAD modellers give a parametrization only of the boundary of geometrical objects as collections of manifolds in three-dimensional space, while the approximation spaces need to be defined on a three-dimensional volume that represents the computational domain. Moreover, these manifolds are often obtained by trimming. In this respect, isogeometric methods have taken the first fundamental steps towards a satisfactory solution and have undoubtedly created strong interest in these questions, promoting interaction between different scientific communities.

3.30 Interactive Deformation of Virtual Paper

Camille Schreck (IST Austria – Klosterneuburg, AT)

Joint work of Camille Schreck, Damien Rohmer, Stefanie Hahmann, Marie-Paule Cani, Shuo Jin, Charlie Wang, Jean-Francis Bloch
Main reference C. Schreck, D. Rohmer, S. Hahmann, M.-P. Cani, S. Jin, C. Wang, J.-F. Bloch, "Nonsmooth Developable Geometry for Interactively Animating Paper Crumpling", ACM Trans. Graph., 35(1):10:1–10:18, ACM, 2015.
URL http://dx.doi.org/10.1145/2829948

Although paper is a very common material in our every-day life, it can hardly be found in 3D virtual environments. Indeed, due to its fibrous structure, paper material exhibits complex deformations and sound behaviour which are hard to reproduce efficiently using standard methods. Most notably, the deforming surface of a sheet of paper is constantly isometric to its 2D pattern, and may be crumpled or torn leading to sharp and fine geometrical features. We propose to combine usual physics-based simulation with new procedural, geometric methods in order to take advantage of prior knowledge to efficiently model a deforming sheet of paper. Our goals are to reproduce a plausible behavior of paper rather than an entirely physically accurate one in order to enable a user to interactively deform and create animation of virtual paper.

3.31 A Framework for Isogeometric Analysis on Unstructured Quadrilateral Meshes

Hendrik Speleers (University of Rome "Tor Vergata", IT)

Joint work of Deepesh Toshniwal, Hendrik Speleers, Thomas J. R. Hughes
Main reference D. Toshniwal, H. Speleers, and T. J. R. Hughes, "Smooth cubic spline spaces on unstructured quadrilateral meshes with particular emphasis on extraordinary points: Geometric design and isogeometric analysis considerations", Computer Methods in Applied Mechanics and Engineering, 327:411–458, Elsevier, 2017.
URL https://doi.org/10.1016/j.cma.2017.06.008

CAD representations of arbitrary genus geometries with a finite number of tensor-product polynomial patches invariably lead to surface representations with unstructured quadrilateral meshes and, thus, extraordinary vertices. Construction of smooth splines over such unstructured quadrilateral meshes is of considerable interest within the field of isogeometric analysis, and a myriad of approaches have been explored that focus on the design and analysis of such geometries.

In this talk we present an alternative approach towards construction of smooth splines over unstructured quadrilateral meshes. Acknowledging the differing requirements posed by design (e.g., the convenience of an intuitive control net) and analysis (e.g., good approximation behavior), we propose the construction of a separate, smooth spline space for each while ensuring isogeometric compatibility. A key ingredient in the approach is the use of singular parameterizations at extraordinary vertices. We demonstrate the versatility of the approach with several applications in design and analysis. The constructed spline spaces show superior approximation behavior, and seem to be well behaved even at the singularities.

3.32 Using Machine Learning Techniques in Geometric Modeling

Georg Umlauf (HTWG Konstanz, DE)

In this talk two applications of machine learning approaches to geometric modeling problems were presented: knot placement for b-spline approximation and shape classification of point clouds.

Selecting knot values to receive good approximation results is a challenging task. Proposed approaches range from parametric averaging to genetic algorithms. We propose the use of Support Vector Machines (SVMs) for finding suitable knot vectors in B-spline curve approximation. The SVMs are trained to distinguish between locations along the curve that are well or not well suited as knots in the parametric domain. This score is based on different geometric features of a parameters corresponding point in the point cloud. A score weighted averaging technique is used to produce the final knot vector. We further propose a method to use the score weighted averaging technique for t-spline surface approximation.

In the reverse engineering process one has to classify parts of point clouds with the correct type of geometric primitive. Features based on different geometric properties like point relations, normals, and curvature information can be used to train SVMs. These geometric features are estimated in the local neighborhood of a point of the point cloud. The multitude of different features makes an in-depth comparison necessary. Instead unsupervised learning methods, e.g. auto-encoders, can be used to learn the relevant features simultaneously with the classification task. These automatic features can be visualized to interpret the learned classification criteria.

3.33 A Multi-sided Bézier Patch with a Simple Control Structure

Tamas Várady (Budapest University of Technology and Economics, HU)

Joint work of Tamas Várady, Peter Salvi

Generalized Bézier (GB) patches represent a new multi-sided scheme; they are compatible with quadrilateral Bézier patches and inherit many of their properties. Boundaries and corresponding cross-derivatives (ribbons) are specified as standard Bézier surfaces of arbitrary degrees. The surface is defined over a convex polygonal domain; local coordinates are computed using generalized barycentric coordinates. The control structure is simple and intuitive, control points are associated with a combination of bi-parametric Bernstein functions, multiplied by rational terms.

The input ribbons may have different degrees, but when the final patch is built, the representation will have a uniform degree. Interior control points – other than those specified by the user – are placed automatically by a special degree elevation algorithm, or can be used for shape optimization or approximation of point clouds. GB patches connect to adjacent Bézier surfaces with G1 or G2 continuity. Several examples will be shown to illustrate the proposed scheme.

References

1 T. Várady, P. Salvi, G. Karikó. *A Multi-sided Bézier Patch with a Simple Control Structure*. Computer Graphics Forum 35 (2), 307–317
2 T. Várady, P. Salvi, I. Kovács, *Enhancement of a multi-sided Bézier surface representation*, Computer Aided Geometric Design 55, 69–83

3.34 Shape Optimization for Fine-scale Structure Design for 3D Printing

Denis Zorin (New York University, US)

Additive fabrication technologies are limited by the types of material they can print: while the technologies are continuously improving, still only a relatively small discrete set of materials can be used in each printed object. At the same time, the low cost of introducing geometric complexity suggests the alternative of controlling the elastic material properties by producing microstructures, which can achieve behaviors significantly differing from the solid printing material.

The presentation focused on results related to design of printable microstructures, including identifying topologies allowing for maximal coverage of effective material properties, printability constraints, stress concentration minimization and worst-case load analysis.

4 Working groups

4.1 Additive Manufacturing Focus Group

Tor Dokken (SINTEF – Oslo, NO), Thomas A. Grandine (The Boeing Company – Seattle, US), and Georg Muntingh (SINTEF – Oslo, NO)

In a discussion session on the interactions between additive manufacturing, computer-aided geometric design, and isogeometric analysis, we identified future challenges and main areas of interests.

The following participants were the most active in the discussion:
- Tor Dokken (SINTEF – Oslo, NO)
- Ron Goldman (Rice University – Houston, US)
- Thomas A. Grandine (The Boeing Company – Seattle, US)
- Jens Gravesen (Technical University of Denmark – Lyngby, DK)
- Bert Jüttler (Universität Linz, AT)
- Géraldine Morin (University of Toulouse, FR)
- Georg Muntingh (SINTEF – Oslo, NO)
- Malcolm A. Sabin (Numerical Geometry Ltd. – Cambridge, GB)
- Giancarlo Sangalli (University of Pavia, IT)
- Timothy L. Strotman (Siemens – Milford, US)

IGA for AM

What will it take, in practice, to incorporate IGA into optimal design of highly engineered products? Although the technology is very promising, there are (as of yet) very few inroads in actual industry practices. Trivariate IGA for AM is a promising direction.

Tools for design and visualisation in AM

In the past there was a focus on developing tools for design and visualisation of B-rep models. The solids now encountered in additive manufacturing require a much richer inner structure, to represent aspects such as densities, multimaterial properties, and microstructures. There is a similar need for tools for design and visualisation of objects with these properties, requiring new techniques and technologies.

Multiscale modelling for AM

Early research is being carried out on printing of structural members for very large objects, such as buildings and jet airplanes. There are major challenges in controlling the process well enough, in making the technology affordable at that scale, and in the design processes. In particular, the scale of the final printed part and the scale of the nanostructure from which it is constructed can span several orders of magnitude, raising the specter of hard mathematical challenges properly handling multiscale modelling. Moreover, multiscale design paradigms are still in their infancy, and much important research remains to be done to understand the most promising and effective technologies for dealing with this.

CAGD for AM

How to make the CAGD representation relevant for AM? In general the physics define the design and the geometry is the fall-out of the design, i.e., form follows function. Hence CAGD will have to focus on the paradigms of the individual AM processes. For instance, trimming is a natural construction in the context of subtractive manufacturing. Is a different paradigm needed for additive manufacturing? An emerging paradigm in AM is multiscale design based on microstructures of objects with a complex topology. The design of such microstructures representing the rich inner structure and properties (density, topology, elastic textures, hardness, etc.) of additively manufactured parts will be a challenge in CAGD.

Functional composition

The tiling of existing design into hexahedral microstructures requires hex-meshing. It is an old trick to generate such meshes by precomposing a spline map representing a trimmed surface by a map from the unit square to the trimmed parameter domain. The same idea has been used for generating porous geometry using tileable microstructures, keeping the number of control points relatively low. An issue is that functional composition causes the degree of the splines to go up, which can be a problem for the analysis. These compositions of spline mappings are interesting in many ways in applications, and the time is ripe to make this part of the CAGD discussion.

4.2 Isogeometric Analysis Focus Group

Tor Dokken (SINTEF – Oslo, NO), Thomas A. Grandine (The Boeing Company – Seattle, US), and Georg Muntingh (SINTEF – Oslo, NO)

Joint work of The participants in the Dagstuhl seminar

Isogeometric analysis (IGA) is a highly promising technology, and yet few industrial inroads have been made. Part of the reason for this is that it represents a disruptive technology, i.e., it requires substantial changes to the industrial flow of digital information in at least two places, both CAD and CAE. Moreover, in the general case it requires thinking about geometric modelling from a volumetric point of view rather than traditional boundary-represented solid modelling. It may prove fruitful to explore boundary element methods and shell methods for structures as a means of getting started, since those types of analysis only require surface models. As with the introduction of other transformative technologies in industry, the path forward will be paved with substantial and successful pilot projects and programs that build customer demand for the technology.

Status of IGA as of 2017

Isogeometric analysis was introduced by Professor Tom Hughes in 2005, and it has not progressed as fast as expected at the time. The community is growing too slow: The first IGA conference in Austin, Texas (2011) attracted 100 participants, while the conference in Pavia, Italy (2017) attracted 187 persons. Success has been reported on IGA for shell formulations, for instance in the LS-Dyna integration. In addition, work on volumetric IGA has started, based on LR B-splines, T-splines and the FEAP-prototype. The Interdisciplinary

Geometric Activities in IGA have brought people from CAD and analysis together, such as in this Dagstuhl seminar.

What is blocking the progress of IGA?

- Early claims that IGA would solve interoperability between Analysis and CAD blocks progress, as interoperability issues still exists – although the use of B-splines in IGA helps.
- Today CAD-systems are based on boundary-representations (shells), and CAD must move to mathematical volumes to properly support volumetric IGA. There is a technology gap between 2-variate CAD and 3-variate analysis. The legacy of existing CAD-models is blocking a move as well.
- In CAD-systems the mathematics is hidden from the user. This makes it difficult to design for IGA.
- There are well-established CAD/FEM-solutions, conservative vendors and customers. Introduction of new technology in systems with many user seats can be painful. Current codes that work well cannot easily be adapted to IGA. The analysis vendors of IGA must be convinced to use IGA.
- For many IGA application areas well-established (industrial) FEM-based solutions exist. So far hundreds of person-years have been invested in IGA research. This competes with thousands of person-years invested in FEM research.
- Most of the hard problems in analysis are vector-valued, whereas IGA has focused on scalar-valued problems. Scalar-interpolants form a tiny niche.
- As of now no industrial funding has been triggered, which is for instance the case for machine learning. Industry wants, in general, fast results.
- IGA has been very good for science. Lots of researchers now understand higher order elements and how to adapt them to analysis frameworks. In practice, there are bottlenecks. For instance, the basis functions necessary for a simple geometry cannot be used to solve Maxwell's equations.
- Most of the work on, for instance, proving error estimates is done theoretically, and it is not used in practice.

How can the progress of IGA be improved?

Killer applications for IGA must be identified, where FEA is outperformed. The following research directions are promising.

- There are (topology) optimization challenges that can only be solved using a tight coupling of geometry and analysis, as provided by IGA.
- Modelling and simulation in medicine, for instance in prosthesis design. Dental implants and hearing aids have a direct connection to additive manufacturing. These require a shift to analysis based on novel, volumetric representations, and IGA can fill this need.
- There have been IGA work programmes where several geometers and analysts have been put together, solving certain equations not previously solved by FEM.

Note that although killer applications are necessary, they are not sufficient. It is not realistic to expect the industry to replace their systems anytime soon, as long as current systems are adequate. The following additional measures can be taken to expedite the industrial uptake of IGA.

- Improve the CAD end of IGA.

- Focus on areas where traditional CAD and FEM are not sufficient, such as additive manufacturing.
- Prepare for IGA in standards, such as STEP ISO 10303.
- Apply to areas beyond engineering. In computer science, for instance, the focus is on discrete geometries, but these are not sufficient for analysis.
- Increase the teaching of geometry, by getting IGA into the computer science, applied maths and engineering curricula, which have traditionally had a focus on discrete geometries.
- Expand the scope of IGA, to increase the total IGA research activity. This can be compared to machine learning, where much of the funding comes from industries. This would also ensure a focus on the practical side, to complement the disproportionate theoretical focus in IGA.
- Start from the models resulting from analysis, and improve these using CAD concepts.

4.3 Big Data Focus Group

Thomas A. Grandine (The Boeing Company – Seattle, US), Tor Dokken (SINTEF – Oslo, NO), and Georg Muntingh (SINTEF – Oslo, NO)

In a discussion session on Big Data in geometric modelling, we identified future challenges and main areas of interests. This includes both the application of Big Data algorithms to modelling and the application of geometric algorithms to Big Data problems. In this report we briefly summarize this discussion.

The following participants were the most active in the discussion:
- Chandrajit Bajaj (University of Texas – Austin, US)
- Heidi Dahl (SINTEF – Oslo, NO)
- Thomas A. Grandine (The Boeing Company – Seattle, US)
- Panagiotis Kaklis (The University of Strathclyde – Glasgow, GB)
- Benjamin Karer (TU Kaiserslautern, DE)
- Georg Muntingh (SINTEF – Oslo, NO)
- Jörg Peters (University of Florida – Gainesville, US)

Design of microstructures

In additive manufacturing the vast freedom of design comes with a high complexity. To deal with this complexity, one strategy is to introduce an intermediate scale, for instance based on cells (e.g. voxels), each of which contains a microstructure. Optimization then becomes feasible, as it can proceed in terms of the parameters of the individual microstructures, to determine a heterogeneous object with optimal properties not attainable by the bulk material. In this sense properly designed microstructures can bridge the gap in complexity. Such multiscale design and manufacturing has the potential to create vast amounts of intermediate data, even though the initial and final design is of limit size.

The digital twin

Another example of big geometric data is the digital twin, which is a concept well known for instance in the aerospace and shipping industry. For each airplane/ship that leaves the factory/ship yard, there is a full-blown virtual model of several terabytes in size. This digital twin is systematically updated with all modifications that happen during its lifecycle, which typically lasts at least 15 years. Such modifications include information on hazardous materials, repairs, etc. This information is often acquired by scanning the internal parts of ships.

Industrial Big Data

Another example of Big Data comes in the form of operational and sensor data, such as pressures and velocities. Although these are not of a geometrical nature per se, they typically have a spatial distribution that should be taken into account. Copious amounts of such data is produced and collected, but how should it be processed? Drawing an analogy to newspapers: will you keep a complete archive, or will you identify clippings? How can we grab and organize the relevant informational content? In certain cases the government demands that collected data remains unprocessed, and there is a sense that one cannot lose any data. More often the goal should be to extract relevant information from compressed data with high accuracy.

Factory models

As factories grow in size, so does their complexity. This complexity can be managed by dynamically building large geometric models of the factory. Such models can track everything that moves, including the parts that are to be assembled or produced. If a change to this flow of parts is proposed, this can be simulated first and checked for viability.

A challenge here is that buildings of large size are themselves subject to changes. For instance, rain can cause the ceiling to move by centimeters, moving and misaligning the sensors, for which millimeter precision is desired.

Earth science and data from space

Data collected by space missions, from the far universe to the solar system and observational science, is another area with a strong interest for manipulating Big Data. Space agencies are also very interested in tracking space debris.

The further our sensors extend, from underground to outer space, the more data is collected. Often the problems encountered have a lot in common, but sometimes the presence of unique features makes new techniques necessary.

Design recognition

How can one automatically recognize past designs for the purpose of future design? A large industrial group has implemented a number of such systems, and within a few years each has fallen into disuse. They keep reimplementing and discarding it. What is preventing the adoption? Is it fundamental? Is it part of the drive to create something new? One reason is that assembling to order works, unless you keep changing the ingredients. Such updates are not local, and everything that was previously in balance has to change. There is a need for machine learning that plays nice with all aspects involved, including geometric structures and compressed sensing.

5 Open problems

5.1 Technical position paper on trivariates

Malcolm A. Sabin (Numerical Geometry Ltd. – Cambridge, GB)

A Brep does not fail to describe a solid. What it fails to do is to provide a boundary-aligned parametrisation of a non-trivial solid or of the cells of a Whitney stratification.

Both for additive manufacture and for IGA of solids we need a parametrisation to provide a 'scaffolding' for the description of internal structure.

Within a Brep modeller there is actually a clean structure. There are topological entities, linked by a BOUNDS relation, such as vertices, edges, faces, cells.

Each of the topological entities is given shape and position by embedding it in a geometrical entity which can be expressed by a LIESIN relation.

These two relations are pretty well all the information the topology needs. The geometry also needs representations of shape, and although a rich modeller may have alternatives, it is possible for NURBS to provide all the embeddings. Except that the solid cells don't need any explicit embedding. The default is $\mathbb{R}^3$, and we can see a cell as being bounded by faces as trimming just like the trimming of a face by edges.

There are problems with trimming, and for many purposes it might be more convenient to have a boundary-aligned parametrisation in both bivariate and trivariate cells. If you want to do this in a way which both

(i) has the trivariate parametrisation compatible at the boundary with the bivariate parametrisations of the faces, and
(ii) avoids serious distortion at the corners of the parametrisation,

we need to split the solid cells into the bricks for which there is a natural trivariate parametrisation. This is exactly the hex meshing problem which Sandia have been addressing for many years.

I observe that

(a) there are extraordinary edges analogous to the exraordinary vertices of a subdivision surface (but I am not assuming subdivision geometry!)
(b) it is not necessary to use valencies other than 3, 4 and 5.
(c) if the valencies are so limited, the number of possible ways in which these extraordinary edges can meet is limited to about a dozen.

The actual geometry can be defined either

(i) by using subdivision,
(ii) by using the analogue of finite-filling quads which still maintain the necessary continuity, or
(iii) by going to the dual where there are ordinary edges, but a dozen or so different 'n-faced cells'.

Of these only the first supports nested spaces of basis functions to meet the full requirements of field analysis. This may be a weak requirement.

Current CAD systems do not handle faces in this way, and so trying to get them to support it for solid cells would be inviting them to put a serious self-inconsistency into the theoretical model which underpins the code. Adding trimmed $\mathbb{R}^3$ would almost certainly be a much easier thing for them to implement.

Another issue is that the tensor product representations that we have are considerably better at representing spatial frequencies which are aligned with the iso-parametric directions than with those that are skew. Using the same basis for both the position field and the under-load displacement field (in the case of elasticity analysis) is making the assumption that the dominant spatial frequencies of the two are well-aligned. Trying to use the same basis for the optimisation displacement fields is, of course, a non-starter because of sheer volume of data.

Participants

- Gudrun Albrecht
Universidad Nacional de Colombia – Medellin, CO
- Marco Attene
CNR – Genova, IT
- Chandrajit Bajaj
University of Texas – Austin, US
- Falai Chen
Univ. of Science & Technology of China – Anhui, CN
- Heidi Dahl
SINTEF – Oslo, NO
- Tor Dokken
SINTEF – Oslo, NO
- Gershon Elber
Technion – Haifa, IL
- Daniela Fußeder
MTU Aero Engines – München, DE
- Carlotta Giannelli
University of Firenze, IT
- Christina Gillmann
TU Kaiserslautern, DE
- Ron Goldman
Rice University – Houston, US
- Thomas A. Grandine
The Boeing Company – Seattle, US
- Jens Gravesen
Technical University of Denmark – Lyngby, DK
- Hans Hagen
TU Kaiserslautern, DE
- Stefanie Hahmann
INRIA Grenoble Rhône-Alpes, FR
- Dianne Hansford
Paradise Valley, US
- Bert Jüttler
Universität Linz, AT
- Panagiotis Kaklis
The University of Strathclyde – Glasgow, GB
- Hongmei Kang
Univ. of Science & Technology of China – Anhui, CN
- Benjamin Karer
TU Kaiserslautern, DE
- Ligang Liu
Univ. of Science & Technology of China – Anhui, CN
- Tom Lyche
University of Oslo, NO
- Carla Manni
University of Rome “Tor Vergata”, IT
- Nicolas Mellado
Paul Sabatier University – Toulouse, FR
- Gëraldine Morin
University of Toulouse, FR
- Georg Muntingh
SINTEF – Oslo, NO
- Jörg Peters
University of Florida – Gainesville, US
- Konrad Polthier
FU Berlin, DE
- Xiaoping Qian
University of Wisconsin – Madison, US
- Ulrich Reif
TU Darmstadt, DE
- Malcolm A. Sabin
Numerical Geometry Ltd. – Cambridge, GB
- Maria Lucia Sampoli
University of Siena, IT
- Giancarlo Sangalli
University of Pavia, IT
- Camille Schreck
IST Austria – Klosterneuburg, AT
- Hendrik Speleers
University of Rome “Tor Vergata”, IT
- Timothy L. Strotman
Siemens – Milford, US
- Georg Umlauf
HTWG Konstanz, DE
- Tamas Várady
Budapest University of Technology and Economics, HU
- Denis Zorin
New York University, US

Report from Dagstuhl Seminar 17222

Robust Performance in Database Query Processing

Edited by

Renata Borovica-Gajic[1], Goetz Graefe[2], and Allison Lee[3]

1 **The University of Melbourne, AU,** renata.borovica@unimelb.edu.au
2 **Google – Madison, US,** goetzg@google.com
3 **Snowflake Computing Inc. – San Mateo, US,** allison.lee@snowflake.net

Abstract

The Dagstuhl Seminar 17222 on "Robust performance in database query processing", held from 28/May until 02/June 2017, brought together researchers from academia and industry to discuss aspects of robustness in database management systems that have not been addressed by the previous instances of the seminar. This article summarizes the main discussion topics, and presents the summary of the outputs of four work groups that discussed: i) updates and database utilities, ii) parallelism, partitioning and skew, iii) dynamic join sequences, and iv) machine learning techniques used to explain unexpected performance observations.

Seminar May 28–2, 2017 – http://www.dagstuhl.de/17222
1998 ACM Subject Classification H.2 Database Management, H.2.2 Physical Design: Access Methods, H.2.4 Systems: Parallel databases, Query processing, Relational databases
Keywords and phrases Robust Query Performance, Database Management Systems, Adaptive Query Processing, Query Optimization, Query Execution, Updates, Parallelism, Data Skew
Digital Object Identifier 10.4230/DagRep.7.5.169

1 Executive Summary

Renata Borovica-Gajic
Goetz Graefe
Allison Lee
Glenn Paulley

The Dagstuhl Seminar 17222 on "Robust performance in database query processing" assembled researchers from industry and academia for the third time to discuss robustness issues in database query performance. The seminar gathered 24 researchers around the world working on plan generation and plan execution in database query processing and in cloud-based massively parallel systems with the purpose to address the open research challenges with respect to the robustness of database management systems.

Delivering robust query performance is well known to be a difficult problem for database management systems. All experienced DBAs and database users are familiar with sudden disruptions in data centers due to poor performance of queries that have performed perfectly well in the past. The goal of the seminar is to discuss the current state-of-the-art, to identify specific research opportunities in order to improve the state-of-affairs in query processing, and to develop new approaches or even solutions for these opportunities.

Unlike the previous seminars, the organizers (Renata Borovica-Gajic, Goetz Graefe and Allison Lee) this time attempted to have a focused subset of topics that the participants discussed and analyzed in more depth. From the proposed topics on algorithm choices, join sequences, updates, database utilities, parallelism and skew, column stores, physical database design, and explainability of non-robust query performance, the participants chose four topics and formed four work groups: i) one discussing updates and database utilities, ii) one discussing parallelism and skew, iii) one discussing join sequences, and iv) one focusing on the explanations of the sources of non-robust performance.

Upon choosing the topics of interest, the organizers then guided the participants to approach the topic through a set of steps: by first considering related work in the area; then introducing metrics and tests that will be used for testing the validity and robustness of the solution; after metrics, the focus was on proposing specific mechanisms for the proposed approaches; and finally the last step focused on the implementation policies.

The seminar thus spent its first day on reviewing prior related work, with a special emphasis on the pieces of work that appeared following the previous instances of the seminar: benchmarks (Dagstuhl 12321 [4, 6, 7]), Smooth Scan [2], and Generalized join [3]. Tuesday was spent on defining metrics and tests. On Wednesday, the participants discussed possible alternative approaches and hiked together in the woods. Thursday was focused on driving one chosen approach to specific mechanisms. Finally, we spent Friday on discussing the policies and presented the overall progress.

At the end of the week, each group was hoping to continue their work towards a research publication. The group on parallelism and skew was hoping to publish first a survey on forms of skew and existing remedies for skew. The work group on dynamic join sequences even had a working prototype by the end of the seminar. The reports of work groups are presented next.

2 Table of Contents

Executive Summary
Renata Borovica-Gajic, Goetz Graefe, Allison Lee, and Glenn Paulley 169

Working groups

Updates and database utilities
Ilia Petrov, Jiaqi Yan, Thanh Do, Knut Stolze, John Cieslewicz, and Goetz Graefe 172

Parallel Join Processing with Skew
Peter A. Boncz, Alfons Kemper, Angelos Anadiotis, J. Christoph Freytag, Kai-Uwe Sattler, and G. N. Paulley 172

Explaining unexpected performance in database query execution
Surajit Chaudhuri, Anisoara Nica, Marcin Zukowski, Fabian Hueske, Immanuel Trummer, Tahir Azim, and Renata Borovica-Gajic 174

Dynamic join sequences
Campbell Fraser, Bart Samwel, Thomas Neumann, Srinivas Karthik, and Allison Lee 175

Summary 178

Participants 180

3 Working groups

3.1 Updates and database utilities

Ilia Petrov, Jiaqi Yan, Thanh Do, Knut Stolze, John Cieslewicz, and Goetz Graefe

Within Dagstuhl Seminar 17222, the work group on updates and database utilities explored mechanisms and policies for database updates including very large updates as common in load utilities. The principal mechanisms for improving performance and scalability focus on avoiding write amplification, e.g., reading and writing an entire database page for an update-in-place modifying just a few bytes, doing the same in multiple indexes and materialized views, logging updates of a few bytes in log records with large headers, etc. Instead of updates-in-place, many systems use variations of very traditional master tapes and change tapes, often called deltas and sometimes forced by append-only file systems. For ordered indexes, the resulting storage structures are log-structured merge trees (really forests with many trees!) and partitioned b-trees (a single b-tree with many internal partitions).

These storage structures may improve not only efficiency and scalability but also robustness of performance, which we understand to mean linear or near-linear (eg N log N) execution costs when expressed as function of size of database or table, size of input or change set, and size of transactions (bulk load vs trickle load). The opposite of robust performance are cost functions with cliffs. Our focus here is on measured performance rather than theoretical performance of algorithms or data structures; we expect that the measured performance correlate very highly with a linear or near-linear regression line.

We hope to capture and extend our results in a publication, e.g., to a workshop on performance benchmarking or on database query processing. Rather than focus on any one proprietary system or on a specific set of techniques, this publication will be about benchmarking and metrics for updates and for database utilities; and rather than focus on single-thread efficiency or multi-node scalability, it will be about measuring robustness of performance.

3.2 Parallel Join Processing with Skew

Peter A. Boncz, Alfons Kemper, Angelos Anadiotis, J. Christoph Freytag, Kai-Uwe Sattler, and G. N. Paulley

The Dagstuhl Workshop 17222 workgroup on Parallelism and Skew explored a range of topics surrounding the ability for a database system to mitigate the problems of data skew in processing join queries in both tightly- and loosely-coupled database systems. While the problems surrounding data skew have been studied over the past thirty years, and some work has documented well some of the fundamental aspects of data skew in join processing, the workgroup felt that an in-depth look at the state of the art was necessary, for a number of reasons. First, much of the work on joins with skewed data was published between 25 and 30 years ago when system performance tradeoffs – for example, the relative speed of disk to main memory – was considerably different than it is today. Second, our workgroup focus was not merely on raw performance but on characterizing robust behaviour, a viewpoint largely missing from prior work. Third, it quickly became desirable to define a robustness benchmark for joins over skewed data so that we could measure performance degradation and compare systems using a common, agreed-upon metric. Fourth, we were aware of some recent research in the literature, notably Flow-Join from this year's ICDE conference in Helsinki and co- authored by two of the members of seminar 17222, that offered advantages over existing, more well-known approaches [10].

During the seminar the workgroup explored three closely-related ideas: a robustness benchmark for join processing in the face of data skew; a survey of existing techniques from the literature; and improvements to current techniques to further mitigate the presence of skew by either avoiding known 'performance cliffs' at query optimization time, or adapting to circumstances at query execution time to avoid such cliffs on one or more threads or processes that will causes overall query latency to increase. During the seminar, the group concentrated on the first two topics. Our primary goal during the seminar was to define a robustness benchmark for join processing in the face of skew, as our feeling was that customers and users have expectations of linear or near-linear (e.g. N lg N) execution costs when expressed as a function of the size of a database or table, and as well expect near-linear performance with respect to increases in memory size and the relative number of CPU cores. The group considered several possible functions, but eventually decided upon a metric R that exponentially degrades as query execution times become less predictable. We then considered options about to construct such a benchmark, and looked at several promising and existing options: the Social Network benchmark, the Star Schema benchmark, TPC-DS, and a modified version of TPC-H. We decided upon the latter option, and discussed the different types of skew that could be placed into the benchmark. We also discussed modifying the benchmark's existing queries, or changing them outright, to determine how best to illustrate the effects of skew.

In the case of skewed column distributions, it makes sense to have different parameter equivalence classes; for example, in the TPC-DS benchmark, these are values from the same step in the step functions that it uses. These are parameters from a skewed distribution, that despite being skewed, still has multiple individual values with the same frequency. The goal is to be able to run a particular query variant with different parameters but mostly identical behavior; that is, values with the same frequency imply similar intermediate result cardinalities.

However, in a skewed distribution, there may be multiple such equivalence classes, e.g. some highly frequent values (all the same frequency) and some low frequency values (all the same frequency). This leads to the concepts of query variants, e.g. Q2 of the benchmark may now have two variants, Q2a and Q2b, where variant 2a binds with the frequent parameter values, and 2b with the infrequent ones. Subsequent to the Workshop our focus group plans to modify the DBGEN utility of the TPC-H benchmark to determine the feasibility of our ideas and their usefulness in deriving a robustness metric for a DBMS [11].

Our second thrust during our meetings was to examine the state-of-the-art in parallel join processing to determine how query processing performance could be improved. We identified a collection of seminal papers from the existing literature as a starting point for our survey. We plan to augment these papers to develop a much more comprehensive survey that we hope will be of value to the research community.

3.3 Explaining unexpected performance in database query execution

Surajit Chaudhuri, Anisoara Nica, Marcin Zukowski, Fabian Hueske, Immanuel Trummer, Tahir Azim, and Renata Borovica-Gajic

Within Dagstuhl Seminar 17222, the work group on "Explaining unexpected performance in database query execution" explored mechanisms for informing users about the source of performance deviation (typically degradation) of a query with respect to prior query executions. The goal of this effort is to increase the user satisfaction and develop trust toward the database management systems by providing explanation for poor query performance. The existence of such a tool would also be highly beneficial for database vendors as a way to mitigate costs of database support, since the tool would reduce the number of angry calls from users requesting justification for a query performance change.

In general, the work group divided the source of deviation into two broad categories: i) expected deviation because of larger query inputs or reduced resources (e.g. lower memory budget), ii) unexpected deviation as a typical consequence of a plan change or a performance cliff in the cost model that triggered higher resource usage (e.g. intermediate results spilling to disk). The first category encompasses "robust", i.e. expected and justified query performance given the query's inputs and outputs, while the second group displays "non-robust" behaviour where the query input is not the reason for observed performance degradation.

Within this particular problem, the group especially focused on a somewhat easier problem of detecting and explaining deviation of a single query template over a sequence of observed query invocations. Templatized queries are typical in modern decision support systems characterized by a frequent repeat of the same plan template with different parameter values (e.g. reports or user-defined functions). Given this context, our problem can be formulated as: Given the past X executions of query Q, detect whether and explain why the invocation X+1 deviates from the expected performance of query Q given the input X+1.

A crucial question to answer is what we consider as "expected" given the new input characteristics and prior query executions. An approach that the group discussed during the seminar is the approach of employing machine learning techniques to train a model of the behavior of the query given its past executions. The group discussed in detail what should constitute a model, which features should be extracted, etc. Alternatively, the group discussed whether techniques such as learning cost models or statistical models could be sufficient to explain the source of performance deviation.

At a high-level, the group settled on an operator-level approach for modeling query performance. For each operator in each query template, the system collects statistics on a variety of metrics, including input and output cardinalities and physical resource usage (CPU, memory, IO). A machine learning model trained for each metric then characterizes its behavior over the query history seen thus far. If a subsequent query is unexpectedly slow, consulting these models provides clues on which metrics were anomalous and could provide a possible explanation for the slowdown. Furthermore, the operator-level approach prevents anomalies from getting compounded or merged up the query tree. Thus, any unexpectedly slow operator can be detected as close to the leaves as possible.

As part of the discussion into this approach, the group reviewed a number of recent papers related to the use of machine learning for modeling and predicting query performance. The goal of this review was two-fold: first, to familiarize the group with the most relevant work in the domain and, second, to check if existing work had previously tackled the problem of explainable query performance. In summary, we found a substantial amount of work on using machine learning for query performance modeling and prediction, but almost no work on explaining the causes when query performance deviates from the trained models.

Malik et al [8] parse out features from the query text of SkyServer queries and combine them with previously seen output cardinalities to train a model for future cardinality estimation. Tzoumas et al [13, 12] use graph theory and machine learning to detect correlations in the input data and thus improve cardinality estimation. Akdere et al [1] demonstrate the effectiveness of machine learning to predict query execution time by using plan-level and operator-level models based on optimizer cost estimates, query parameters and actual runtime statistics. Finally, Wu et al [5] avoid machine learning approaches and instead use additional statistics to predict query execution time by building on PostgreSQL's optimizer cost models. They use offline profiling to estimate the unit cost of reading pages and tuples using different access paths. They then use online sampling for each query to estimate how many pages and tuples will be read by a query. In a series of follow-up papers [14, 15], they incorporate more sophisticated statistical models to improve query performance prediction in the presence of uncertainty and concurrency. Microsoft's Adaptive Query Processing feature [9] adds the ability to test a query's performance with different plans and visually inspect operators for estimated and actual statistics along with possible performance warnings. However, it does not correlate the current query performance with what the user has seen in the past. While each of these techniques improves the state-of-the-art in query performance estimation, the task of explaining why many queries deviate from the estimates remains out of their scope.

The group also gave considerable attention towards proposing benchmarks that will model realistic use cases, while at the same time would catch all sources of "unexpected" performance behavior. As a first step to this goal the work group proposed a benchmark based on significant data correlation and workload skew. As a future step, we plan to consider the impact of resources as well and incorporate them into the test benchmark.

The overall work group plan is to pursue a publication on this topic followed by a supporting software tool whose main focus would be on explaining unexpected performance to the user either graphically or in the form of natural language text.

3.4 Dynamic join sequences

Campbell Fraser, Bart Samwel, Thomas Neumann, Srinivas Karthik, and Allison Lee

The workgroup on operator sequencing in Dagstuhl Seminar 17222 explored techniques to avoid catastrophically bad performance due to poor join order selection by the query optimizer. In a conventional SQL Query Processing engine, the Query Optimizer generates a single execution plan for each query. For a host of reasons, this often results in query performance that is orders of magnitude worse than that of the fastest possible performance for the query. The models used to estimate the cardinalities at intermediate stages of query execution are often too simple to provide accurate estimates, and the statistics on which they rely can be out of date. Parameterised queries are particularly vulnerable because, even with perfect statistics and a rich estimation model, it may be true that no single plan is capable of providing acceptable performance for all possible parameters, for a given query.

The workgroup reviewed existing adaptive execution techniques that have been proposed to overcome this architectural limitation, including re-optimizing fragments of the plan during execution, tuple routing, dynamically rearranging the order of joins, and parameter sensitive plans. Based on previous work and initial ideas for solutions, we generated a list of limitations to the scope of solutions, that we could use to evaluate the usefulness of a solution for a given use case. For instance, previous work on tuple routing was limited to certain join methods (e.g. symmetric hash join), and work on dynamically rearranging the order of joins was limited to linear plans of indexed nested-loops joins.

The group proposed a simple test query and data set that could be used to evaluate the performance of proposed solutions.

```
SELECT *
FROM nations n
     join customers c using (n_id)
     join orders o using (c_id)
     join suppliers s on s.n_id = c.n_id and s.s_id = o_sid
WHERE n.name = :n
```

The query contains a join which can substantially blow up the size of intermediate results, depending on the value of a query parameter (: n). For most (less frequent) parameter values, any join order performs adequately, while for more frequent values, the wrong join order will result in performance that is orders of magnitude worse than other join orders. Throughout our discussions, this query was used to evaluate the potential of various ideas that were proposed. We also proposed a more general technique to evaluate a workload against a proposed solution, which involved generating a large set of random join orders for each query, detecting the "cliff" in the distribution of runtimes of those join orders, and then evaluating the impact of a solution on join orders falling on either side of the cliff.

After brainstorming several possible solutions, we focused on a novel tuple-routing approach. The bookkeeping for the routing is simple, does not increase in size or complexity with the number of joins in the query, and requires only inexpensive telemetry that is already gathered by some query processing engines. An intuitive system of back pressure and coordination is used to route tuples away from inefficient join orders and through efficient orders. The method is primarily aimed at linear indexed nested-loop joins but can also work with bushy plans and with hash tables functioning as "indexes on the fly". Unlike previous adaptive techniques, this method is not index restrictive; a different index can be used for the same table when it appears in different join orders. The algorithm is always making forward progress; it never backtracks or throws away partial results. Figures 1a and 1b show two plan options for our simple test query above. For a frequently occurring value of nation (e.g. 'US'), the plan in Figure 1a performs well, while the plan in Figure 1b performs very badly. Figure 2 shows a tuple-routing plan for this query, with the "good" and "bad" paths annotated. Back pressure will cause most tuples to be routed through the left-most path.

During the seminar, we implemented a prototype inside of a teaching database, which handled linear indexed nested-loops joins only. We compared performance (measured as the total number of tuples produced as intermediate results) of all possible join orders for our test query against our tuple routing scheme. As expected, we found a worst case 2X degradation in intermediate tuples produced versus a good plan, and orders of magnitude improvement in the worst plans.

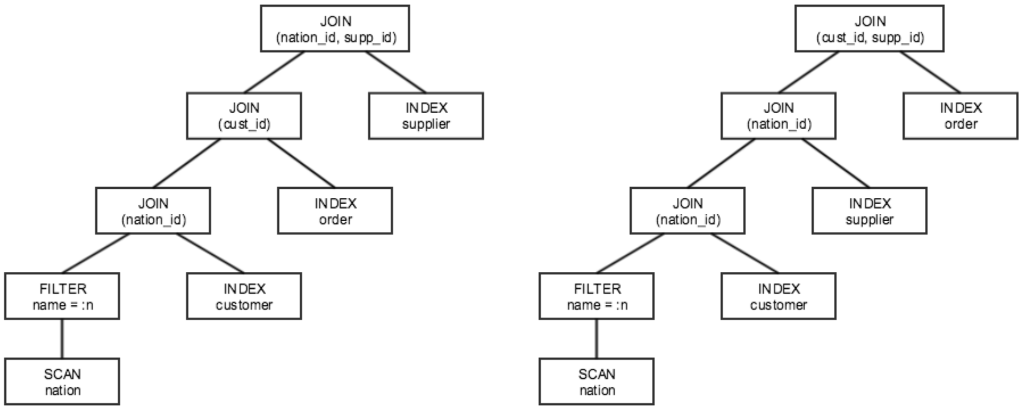

(a) Good plan **(b)** Bad plan

Figure 1 An example of a good and bad plan

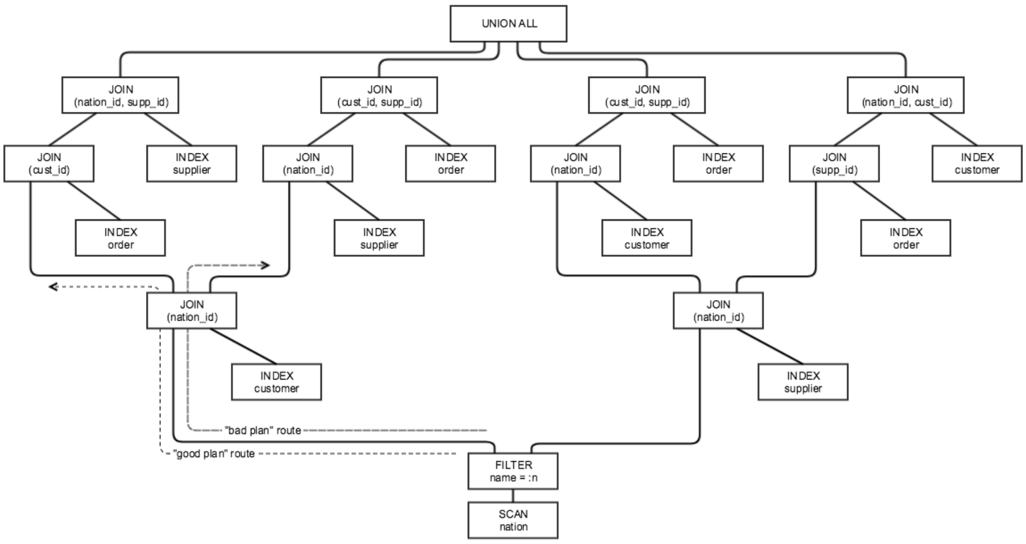

Figure 2 A tuple routing plan

There are several open issues which the group plans to continue work on after the seminar, including:

- policies for when to start to route tuples, to decrease the performance degradation when the optimizer chooses a good plan, and when to stop or decrease tuple routing once a good plan has been found
- implementing back-pressure for non-linear plans,
- allowing hash joins in the plan,
- techniques for swapping the driving table,
- techniques for choosing a subset of join orders in the plan,
- prototyping the solution in a real database system.

4 Summary

State-of-the-art approaches in query processing are oriented toward achieving high performance while robustness of these approaches is often neglected. This may result in queries whose run time fluctuates severely as a result of marginal changes in the underlying data. When talking about robust performance, stable and expected or close to expected performance has to be provided every single time, even in the presence of changes in the selectivity of operators or changes in the system state between compile time and run time. This report presents four approaches toward improving the robustness of database management systems by looking at the impact of updates, parallelism and skew, join orders and explainability of non-robust query performance.

References

1 Mert Akdere, Ugur Çetintemel, Matteo Riondato, Eli Upfal, and Stanley B. Zdonik. Learning-based query performance modeling and prediction. pages 390–401, 2012.
2 Renata Borovica-Gajic, Stratos Idreos, Anastasia Ailamaki, Marcin Zukowski, and Campbell Fraser. Smooth Scan: Statistics-oblivious access paths. In *ICDE*, 2015.
3 Goetz Graefe. New algorithms for join and grouping operations. *Computer Science - R&D*, 27(1):3–27, 2012.
4 Goetz Graefe, Wey Guy, Harumi A. Kuno, and Glenn N. Paulley. Robust query processing (dagstuhl seminar 12321). *Dagstuhl Reports*, 2(8):1–15, 2012.
5 Hakan Hacigumus, Yun Chi, Wentao Wu, Shenghuo Zhu, Junichi Tatemura, and Jeffrey F. Naughton. Predicting query execution time: Are optimizer cost models really unusable? pages 1081–1092, 2013.
6 Martin L. Kersten, Alfons Kemper, Volker Markl, Anisoara Nica, Meikel Poess, and Kai-Uwe Sattler. Tractor pulling on data warehouses. In *Proceedings of the Fourth International Workshop on Testing Database Systems*, DBTest '11, pages 7:1–7:6, New York, NY, USA, 2011. ACM.
7 Martin L. Kersten, Alfons Kemper, Volker Markl, Anisoara Nica, Meikel Poess, and Kai-Uwe Sattler. Tractor pulling on data warehouses. In *Proceedings of the Fourth International Workshop on Testing Database Systems*, DBTest '11, pages 7:1–7:6, New York, NY, USA, 2011. ACM.
8 Tanu Malik, Randal C. Burns, and Nitesh V. Chawla. A black-box approach to query cardinality estimation. pages 56–67, 2007.
9 Microsoft. Microsoft adaptive query processing and diagnostics. `https://www.youtube.com/watch?v=szTmo6rTUjM`.

10 Wolf Rödiger, Sam Idicula, Alfons Kemper, and Thomas Neumann. Flow-join: Adaptive skew handling for distributed joins over high-speed networks. In *ICDE*, pages 1194–1205, 2016.
11 TPC. Tpc-h benchmark. `http://www.tpc.org/tpch/`.
12 Kostas Tzoumas, Amol Deshpande, and Christian S. Jensen. Lightweight graphical models for selectivity estimation without independence assumptions. *PVLDB*, page 2011, 2011.
13 Kostas Tzoumas, Amol Deshpande, and Christian S. Jensen. Efficiently adapting graphical models for selectivity estimation. *The VLDB Journal*, 22(1):3–27, 2013.
14 Wentao Wu, Yun Chi, Hakan Hacígümüş, and Jeffrey F. Naughton. Towards predicting query execution time for concurrent and dynamic database workloads. *PVLDB*, 6(10):925–936, 2013.
15 Wentao Wu, Xi Wu, Hakan Hacigümüş, and Jeffrey F. Naughton. Uncertainty aware query execution time prediction. *PVLDB*, 7(14):1857–1868, 2014.

Participants

- Angelos-Christos Anadiotis
EPFL – Lausanne, CH
- Tahir Azim
EPFL – Lausanne, CH
- Peter A. Boncz
CWI – Amsterdam, NL
- Renata Borovica-Gajic
The University of Melbourne, AU
- Surajit Chaudhuri
Microsoft Research – Redmond, US
- John Cieslewicz
Google Mountain View, US
- Thanh Do
Google – Madison, US
- Campbell Fraser
Google – Kirkland, US
- Johann-Christoph Freytag
HU Berlin, DE
- Goetz Graefe
Google – Madison, US
- Fabian Hüske
data Artisans – Berlin, DE
- Alfons Kemper
TU München, DE
- Allison Lee
Snowflake Computing Inc. – San Mateo, US
- Thomas Neumann
TU München, DE
- Anisoara Nica
SAP SE – Waterloo, CA
- Glenn Paulley
SAP SE – Waterloo, CA
- Ilia Petrov
Hochschule Reutlingen, DE
- Bart Samwel
Google – Amsterdam, NL
- Kai-Uwe Sattler
TU Ilmenau, DE
- Knut Stolze
IBM Deutschland – Böblingen, DE
- Immanuel Trummer
Cornell University – Ithaca, US
- Srinivas Karthik Venkatesh
Indian Institute of Science – Bangalore, IN
- Jiaqi Yan
Snowflake Computing Inc. – San Mateo, US
- Marcin Zukowski
Snowflake Computing Inc. – San Mateo, US