Group Prof. Roger Wattenhofer
Distributed Computing

Professor
Roger Wattenhofer

Mission

We are interested in both theory and practice of computer science and information technology. In our group we cultivate a large breadth of areas, reflecting our different backgrounds in computer science, mathematics, and electrical engineering. This gives us a unique blend of basic and applied research, proving mathematical theorems on the one hand, and building practical systems on the other.

We currently study the following topics: Distributed computing (computability, locality, complexity), distributed systems (Bitcoin), wireline networks (software defined networks), wireless networks (media access theory and practice), social networks (influence), algorithms (online algorithms, game theory), learning theory (recommendation theory and practice). We regularly publish in different communities: distributed computing (e.g. PODC, SPAA, DISC), networking (e.g. SIGCOMM, MobiCom, SenSys), theory (e.g. STOC, FOCS, SODA, ICALP), and from time to time at random in areas such as machine learning or human computer interaction.

Members of our group have won several best paper awards at top conferences such as PODC, SPAA, DISC, MobiCom, or P2P. Roger Wattenhofer has won the Prize for Innovations in Distributed Computing in 2012, for "extensive contributions to the study of distributed approximation". Some projects turned into startup companies, e.g. Wuala, StreamForge, BitSplitters. Several projects have been covered by popular media and blogs, e.g. Gizmodo, Lifehacker, New York Times, NZZ, PC World Magazine, Red Herring, or Technology Review. Some of the software developed by our students is very popular: The music application Jukefox and the peer-to-peer client BitThief have together more than 1 million downloads. A branch of the United States FBI has requested to use a version of BitThief as a tool to uncover illegal activities. About half of the former PhD students are in academic positions, some others founded startup companies.

Curriculum Vitae

Prof. Roger P. Wattenhofer
Professor of Distributed Computing
Head of Computer Engineering and Networks Laboratory

8/2008-present: Full Professor, Distributed Computing Group, Computer Engineering and Networks Laboratory, Department of Information Technology and Electrical Engineering, ETH Zurich
7/2012-8/2013: Sabbatical, Systems and Networking Group, Microsoft Research, Redmond, WA
7/2004-7/2008: Associate Professor, Distributed Computing Group, Computer Engineering and Networks Laboratory, Department of Information Technology and Electrical Engineering, ETH Zurich
10/2006-3/2007: Sabbatical, Macquarie University, Sydney, Australia
10/2001-6/2004: Assistant Professor, Distributed Computing Group, Institute for Pervasive Computing, Department of Computer Science, ETH Zurich, Switzerland
4/2000-10/2001: Post-Doc Researcher, Systems and Networking Group, Microsoft Research, Redmond, WA
4/1999-4/2000: Post-Doc Researcher, Computer Science Department, Brown University, Providence, RI
1995-1999: Ph.D. in Computer Science, Research and Teaching Assistant, Computer Science Department, ETH Zurich. (Advisor Peter Widmayer, Co-Examiner Maurice Herlihy, Additional Expert Nir Shavit)
1990-1995: Studies in Computer Science, ETH Zurich, Switzerland, with Minor in Operations Research

Google Scholar lists more than 200 papers co-authored by Roger Wattenhofer, more than 100 in the reviewing period. These publications have a total of more than 14,000 citations (the majority of them in the reviewing period), and an h-index above 60 (more than 50 in the reviewing period). The Palsberg list of highly cited computer scientists ranks Roger Wattenhofer in the Top 200. Roger Wattenhofer is married, with 3 children.

Group Structure and Resources

Prof. Roger Wattenhofer's group

Our administration is shared with the other groups in the institute. Our research group has a flat organization, all members are PhD students:
Tobias Langner, MSc in CS, University Freiburg, joined in March 2010, currently working on biologically inspired algorithms in multi-agent systems.
Barbara Keller, MSc in CS, ETH Zurich, joined in October 2010, currently working on social networks, in particular how nodes influence each other in theory and practice.
Jochen Seidel, Dipl.-Inform., KIT Karlsruhe, joined in June 2011, currently working on computability in anonymous networks and DNA strand displacement.
Jara Uitto, MSc in CS, University of Helsinki, joined in September 2011, currently working on weak models of computation and recommendation systems.
Philipp Brandes, MSc in CS, University Paderborn, joined in October 2011, currently working on designing mechanisms in SDNs to explore how they can deal with selfish streams.
Klaus-Tycho Förster, Dipl.-Math. and Dipl.-Inform., TU Braunschweig, joined in October 2011, currently working on hardness results and efficient algorithms for updating SDNs.
Christian Decker, MSc in CS, ETH Zurich, joined in March 2012, currently working on analyzing the networking aspects of Bitcoin, a decentralized global currency system.
Michael König, MSc in CS, ETH Zurich, joined in May 2012, currently working on improving the capacity of wireless networks by simultaneous sending yet avoiding interference.
Pascal Bissig, MSc in EE, ETH Zurich, joined in June 2012, currently working on mobile sensing in various applications such as sports or social interaction statistics.
Sebastian Brandt, Dipl.-Math., University Bonn, joined in January 2014, currently working on fast algorithms for finding reasonably small vertex separators in graphs.
Stephan Kollmann, MSc in CS, ETH Zurich, joined in March 2014, currently working on algorithms for approximating sparse user-movie-rating-matrices with provable performance.

Key Figures

	2013	2012	2011	2010	2009	2008
PhD students	10	12	13	10	13	12
Postdocs	-	1	1	0.5	-	-
Senior Scientists	-	-	-	-	-	-
Adm. / Technical / IT Staff	1	1	1	1	1	1
Total	**11**	**14**	**15**	**11.5**	**14**	**13**

Finished PhD theses over the last 6 years

2013: Samuel Welten: *Sensing with Smartphones: Light Authentication, Heavy Personalization and Medical Applications*
Stephan Holzer: *Distance Computation, Information Dissemination, and Wireless Capacity in Networks*
2012: Raphael Eidenbenz: *Coping with Selfishness in Distributed Systems: Mechanism Design in Multi-Core and Peer-to-Peer Systems*
2011: Johannes Schneider: *Decentralized Coordination: Methods and Applications*
Philipp Sommer: *Wireless Embedded Systems: Time, Location, and Applications*
Remo Meier: *Toward Structured and Time-Constraint Content Delivery Systems*
Christoph Lenzen: *Synchronization and Symmetry Breaking in Distributed Systems**
2010: Michael Kuhn: *Understanding and Organizing User Generated Data: Methods and Applications*
Nicolas Burri: *Ultra-Low Power Sensor Networks: Development Tools, Design, and Implementation*
2009: Roland Flury: *Routing on the Geometry of Wireless Ad Hoc Networks*
Olga Goussevskaia: *Computational Complexity and Scheduling Algorithms for Wireless Networks*
Yvonne Anne Pignolet: *Algorithmic Challenges in Wireless Networks: Interference, Energy, and Incentives*
Thomas Locher: *Foundations of Aggregation and Synchronization in Distributed Systems**
2008: Pascal von Rickenbach: *Energy-Efficient Data Gathering in Sensor Networks*
Stefan Schmid: *Dynamics and Cooperation: Algorithmic Challenges in Peer-to-Peer Computing*

The theses marked with an asterisk (*) have been awarded with an ETH Medal.

Research Activities and Achievements

In the following, we discuss a few examples of our ongoing research projects in more detail.

One of our main theoretical interests is algorithmic theory that does not follow the traditional input/output model of computation. One may call this "physical algorithms", algorithms that live in networked systems of active agents. As many physical systems (cars, financial agents, animals, brain cells, you name it) show "algorithmic" behavior, we would like to understand the fundamentals of such networked systems. Let us give a few examples.

Distributed Complexity: What can be computed, and how efficiently, are probably the core questions of computer science. Not surprisingly, in distributed systems and networking research, a core question is what can be computed in a distributed fashion, in a network. More precisely, if nodes of a network must base their decision on their local neighborhood only, how well can they compute or approximate a global optimization problem? Throughout the years, we studied different aspects and problems of locality in great detail – today this body of work has developed into what is known as distributed complexity theory. In 2012, Roger Wattenhofer received the Prize for Innovation in Distributed Computing for this line of work.

Clock Synchronization: Networks often need a common notion of time; consequently clock synchronization in networks seems to be such a fundamental and practically important question that it should have been solved a long time ago. Surprisingly, this is not the case. In a series of papers we studied the theory of clock synchronization, proving the surprising result that two neighboring nodes cannot synchronize their clocks arbitrarily well; indeed one can show that even the best possible protocol will produce a clock skew between neighbors that scales with the logarithm of the network size. This result is tight, as we discovered an algorithm that achieves this lower bound. We also looked into the practical side of clock synchronization, and developed protocols that beat the state of the art in sensor networks considerably.

Wireless Algorithms: Despite the omnipresence of wireless networks, surprisingly little is known about their computational complexity and efficiency. We developed techniques to understand the fundamental communication limits of arbitrary wireless networks, in a reasonable physical model such as SINR. We published several papers tackling this question. These papers were the seed of a new community, now with a yearly workshop called WRAWN, dedicated to this subject.

Software Defined Networks (SDNs): Large providers such as Microsoft, Amazon, or Google operate their own wide area networks that cost them hundreds of millions of dollars per year, yet even their busier links are only utilized 50% on average. This gives rise to SDNs, where the data and control plane are separated, allowing a controller to update network rules to enhance performance. Scheduling flows and updating routing tables under inherent asynchrony and hardware constraints lie at the core of this emerging technology. Starting in 2013, we proved various hardness results and designed efficient update schematics, which were tested in collaboration with industry partners.

Biological Algorithms: "They operate without any central control. Their collective behavior arises from local interactions." This is the mantra of our main research area, distributed computing, however, in this quote, "they" are not nodes in a distributed system. Rather, the quote is taken from a biology paper that studies social insect colonies. Understanding the behavior of insects is interesting both from a biological and a computational perspective. In recent years, we studied various aspects of such systems and focused our efforts particularly on developing mathematical models that allowed us to make both qualitative and quantitative statements on the capabilities of insects. We showed, for example, that ants can effectively collaborate to locate a food source in an infinite grid in nearly optimal time even if each ant only possesses a constant amount of memory.

Bitcoin: Money facilitates the exchange of goods and services (and organizing debt, more honestly). Bitcoin is an attempt to revolutionize the concept of money, using computational methods. While Bitcoin is already operational and is used to buy just about everything, from alpaca socks to flights into space, it still faces some major challenges: How can it scale to handle truly global adoption? How can we enforce security? How can we improve on the overall experience? We are at the forefront of research regarding the networking aspects of Bitcoin.

Social Networks: Social interaction is a big part in everyone's life. Facebook and similar services collect big data sets of digital interactions and relationships between people. We analyze such networks and model the evolution of networks such as co-authorship in publications. This helps to understand social phenomena such as the glass ceiling effect and gives insights in possible counter measures. We also study purely graph theoretic notions of social interactions, e.g. the complexity of social influence.

Recommendation: The goal of a recommendation system is to provide users with an efficient way to find items (books, movies, songs, apps) that they will like. The system can learn the preferences by asking the users about their preferences – the challenge is to minimize the number of questions. We study recommendation systems with the goal to find at least a single good item per user. We designed an algorithm that is nearly optimal and performs well even when compared to an algorithm that knows a probability distribution on the preferences of the users.

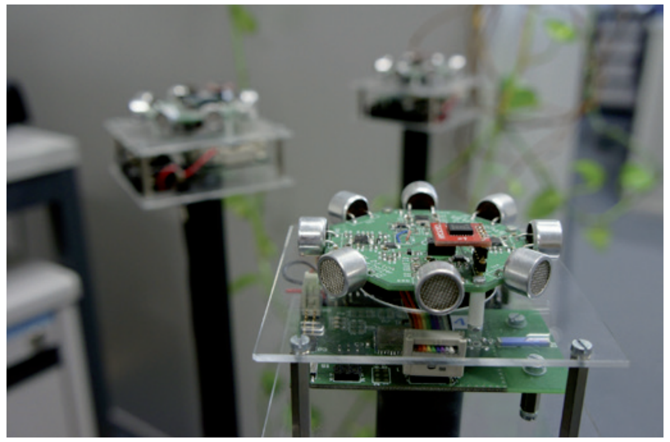

Our SpiderBat platform allows highly accurate distance and angle measurements using ultrasound

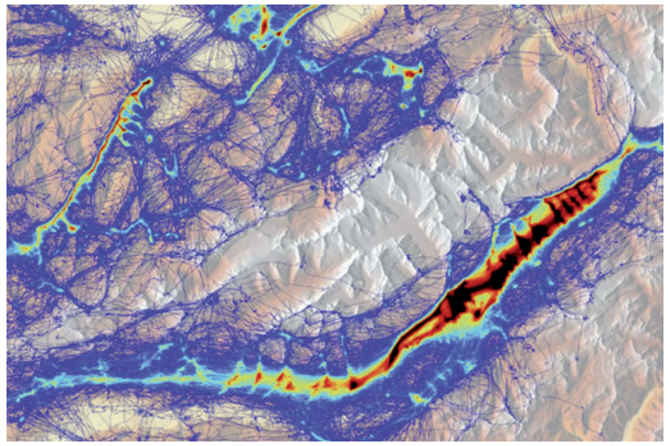

Screenshot of paragliding traces over Swiss alps, from a student project

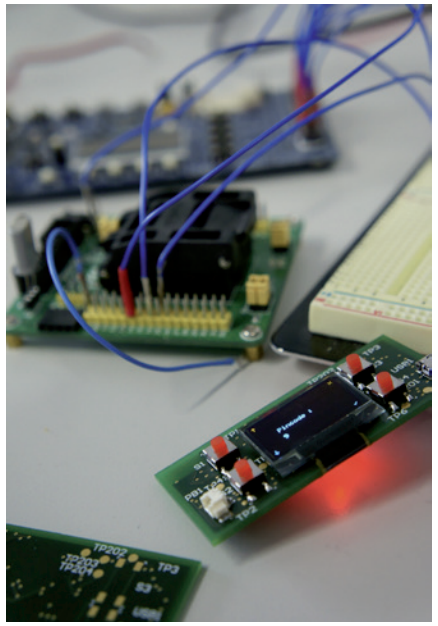

Secure hardware wallet for Bitcoin, designed in a student project

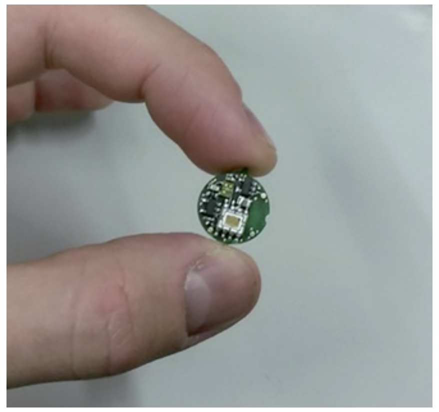

Powerless UV radiation sensor by BitSplitters

Screenshot of an application that allows to quickly map signal strength distributions indoors, developed by students

Selected Publications

Ch. Lenzen et al.: *PulseSync: An Efficient and Scalable Clock Synchronization Protocol,* IEEE/ACM Transactions on Networking, 2014

Ch. Decker, R. Wattenhofer: *Information Propagation in the Bitcoin Network,* 13th IEEE International Conference on Peer-to-Peer Computing (P2P), 2013 (Best paper award)

C.-Y. Hong et al.: *Achieving High Utilization with Software-Driven WAN,* Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM), 2013

H. H. Liu et al.: *zUpdate: Updating Data Center Networks with Zero Loss,* Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM), 2013

Y. Emek, R. Wattenhofer: *Stone Age Distributed Computing,* 31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC), 2013

A. Das Sarma et al.: *Distributed Verification and Hardness of Distributed Approximation.* SIAM Journal on Computing, 2012 (Special issue of selected papers at STOC 2011)

S. Frischknecht et al.: *Networks Cannot Compute Their Diameter in Sublinear Time,* 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA), 2012

Ch. Lenzen, R. Wattenhofer: *Tight Bounds for Parallel Randomized Load Balancing,* 43rd Symposium on Theory of Computing (STOC), 2011

Ch. Lenzen et al.: *Tight Bounds for Clock Synchronization,* Journal of the ACM (JACM), Volume 57, Issue 2, 2010. (Journal version of a FOCS 2008 paper, and a PODC 2009 paper winning the best paper award)

F. Kuhn et al.: *Tight Bounds for Distributed Selection,* 19th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2007 (Best paper award, a popular variant of this paper has been invited by the Communications of the ACM (CACM) magazine, published in the Section Research Highlights, Volume 51, Issue 9, 93-99, Sept. 2008)

Teaching Activities

Roger Wattenhofer currently teaches three courses, plus a seminar. Distributed Systems is a Computer Science 3rd year core course, taught together with Friedemann Mattern from the Computer Science department, with about 100 students. In our part of the course, we primarily discuss fault-tolerance issues (models, consensus, agreement) as well as replication issues (primary copy, 2PC, 3PC, Paxos, quorum systems, distributed storage) and problems with asynchronous multiprocessing (shared memory, spin locks, concurrency).

Discrete Event Systems is an Electrical Engineering 3rd year selective course, with about 50 students. We start out the course by studying popular models of discrete event systems, such as automata, languages, or Petri nets. In the second part of the course, we analyze discrete event systems. We first examine discrete event systems from, an average-case perspective: we model discrete events as stochastic processes, and then apply Markov chains and queuing theory for an understanding of the typical behavior of a system. In the last part of the course we analyze discrete event systems from a worst-case perspective, using the theory of online algorithms, adversarial queuing, and network calculus.

Principles of Distributed Computing is an advanced graduate level course, taught to students from both Computer Science and Electrical Engineering, with about 120 students. Distributed computing is essential in modern computing and communication systems. Examples are on the one hand large-scale networks such as the Internet, and on the other hand multiprocessors such as multi-core laptops. This course introduces the principles of distributed computing, emphasizing the fundamental issues underlying the design of distributed systems and networks: communication, coordination, fault-tolerance, locality, parallelism, self-organization, symmetry breaking, synchronization, uncertainty. We explore essential algorithmic ideas and lower bound techniques, basically the "pearls" of distributed computing.

In addition, we currently supervise about 20 student projects each semester. These projects usually are about the areas of our expertise, however, sometimes we also do some highly speculative and/or fun projects with students. Pictures of these fun projects illustrate this brochure.

This aerial picture is taken from a phone thrown into the air, a wild student project

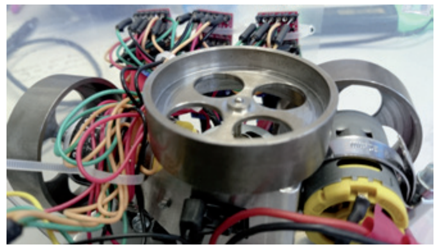

Another student project reduced the blur of these aerial pictures by absorbing the angular momentum with reaction wheels

Professorships | Professor Roger Wattenhofer, Distributed Computing

Service and Transfer

In the reviewing period, Roger Wattenhofer has served as program committee chair or general chair of several renowned international conferences, such as P2P 2014, ICALP 2012, ICDCN 2012, SSS 2011, ICDCN 2009, IPTPS 2007, and PODC 2007. In addition, he was the Committee Chair of the Edsger W. Dijkstra Prize in Distributed Computing in 2008. Moreover, he is a steering committee member of numerous conferences and workshops such as PODC, WRAWN, ICDCN, TAPAS, FOWANC, and a council member of EATCS, the European Association for Theoretical Computer Science. He also was a local organizer of PODC 2010, the annual flagship conference of the Distributed Computing community, for the first time held outside North America, and a co-organizer of several seminars at Schloss Dagstuhl. In addition, during the review period, he was a member of the program committee of 30 international conferences and workshops, and an external examiner on more than a dozen election committees, tenure cases, and PhD defenses.

Among the services at ETH Zurich are: Head of the Computer Engineering and Networks Laboratory, 2008-2009, and again since 2012. Member of the Commission for Studies ("Unterrichtskommission") at the ITET Department, since 2011. Chair of the Commission of Computing Projects (KIM) at the ITET Department, 2009-2012. Member of the Admission Board for Graduate Studies at the ITET Department, since 2006. Representing the ITET Department to prospective students, 2004-2010. Moreover, member of several Election Committees.

A bit more than half of the PhD graduates joined a university or research lab after their PhD: IBM Research (3 PhD students), Microsoft Research (2), ABB Research (2), CSIRO Australia, Hebrew University, MIT, and TU Munich.

Some PhD graduates founded a startup company: 3 former PhD students founded StreamForge, a startup company in the area of video and audio streaming, and one former PhD student founded BitSplitters, a startup company developing a wearable gadget that measures UV radiation. Wuala, a startup company working on secure storage with more than a dozen engineers, was founded by former Master students, and has recently been acquired by Seagate.

Finally, some PhD graduates joined a company: Google (3), AppTornado, Ergon. Our only PostDoc Yuval Emek joined the Technion.

Outlook

We do not plan to change our formula for research and teaching drastically. Our research foci of the late 1990s and early 2000s, sensor networks and peer-to-peer systems, are less important now. New topics have taken their place, and they are thriving as well. Our breadth of research topics allows us to react to trends and developments quickly. However, in the past we have taken pride in the fact that we often started trends, rather than following them.

Regarding teaching, we always have and will continue to change the courses we teach. Our lecture material is available online, and some colleagues around the world are teaching the courses developed by us at their university. In particular the Principles of Distributed Computing course is currently taught at about 20 universities around the world. And even the Ad Hoc and Sensor Networks course (abandoned several years ago from our curriculum) is still taught at a few universities around the world.

Kids are fascinated by the sheer size of the balloon carrying the autonomous weather station, developed in a student project