

## Aapo Kyrölä - Curriculum Vitae

---

*Updated: July 7, 2016*

### CONTACT

#### INFORMATION

Email: akyrola@cs.cmu.edu, akyrola@fb.com

Mobile: +1 650 575 7789

Home Address: 620 Palo Alto Ave, Mountain View, CA, USA

### VERY BRIEFLY

I have founded two pioneering startups and earned my Ph.D. in Computer Science from the Carnegie Mellon University in 2014. My research papers have had large impact and received hundreds of citations in a short time. I have built several challenging and groundbreaking systems from scratch and especially enjoy when I can combine my expertise in systems building with my interests in machine learning, data analysis and algorithms for handling large data. Currently I work at Facebook Inc, which acquired my startup in 2014.

### BRIEFLY

I was founder of one of the biggest virtual world services **Habbo Hotel** (founded year 2000, still going strong) with very strong experience in developing distributed systems with huge number of concurrent users. I hired the initial development team and led it for the first five years. I am particularly proud that many of the engineers I hired have since become successful startup founders on their own right.

I started my Ph.D. studies at the Carnegie Mellon University in Fall 2009, and I am focusing on large-scale machine learning and graph computation. I am advised by professors Guy Blelloch and Carlos Guestrin. My main research project was **GraphChi**, a disk-based graph computation system that can process as big problems as current distributed systems very efficiently, but on just a laptop. GraphChi was covered by the MIT Technology Review in July 2012<sup>1</sup>. **GraphChi-DB** takes the ideas of GraphChi further and is an extremely efficient graph database with state-of-the-art performance (on-going work).

During my Ph.D. studies I founded ProtoGeo, company developing the automatic fully software based activity tracker **Moves** for smartphones. I developed the first prototype of the full technology stack and helped assemble the founding team and raise initial funding for the company. I had significant role in solving the technology problems of making Moves the most intelligent activity tracker in the market. Moves was selected as top app of year 2013 by both Apple and Google.

---

<sup>1</sup> <http://www.technologyreview.com/news/428497/your-laptop-can-now-analyze-big-data/>.

## PUBLICATIONS

Ph.D. dissertation: Aapo Kyrola *Large-scale Graph Computation on Just a PC*, CMU-CS-14-118

Aapo Kyrola, Carlos Guestrin *GraphChi-DB: Simple Design for a Scalable Graph Database – On Just a PC* (arxiv: <http://arxiv.org/abs/1403.0701>)

Aapo Kyrola, Julian Shun, Guy Blelloc, *Beyond Synchronous: : New Techniques for External-Memory Graph Connectivity and Minimum Spanning Forest*, Symposium on Experimental Algorithm, Copenhagen, July 2014.

Aapo Kyrola *DrunkardMob: Billions of Random Walks on Just a PC* ACM Recommender Systems conference 2013.

Aapo Kyrola, G. Blelloch, C. Guestrin. *GraphChi: Large-scale Graph Computation on Just a PC*. OSDI 2012.

Y. Low, J. Gonzalez, Aapo Kyrola, D. Bickson, C. Guestrin, and J.M. Hellerstein. *Distributed GraphLab: A Framework for Machine Learning in the Cloud*. PVLDB 5(8):716-727, 2012.

R. Cheng, J. Hong, Aapo Kyrola, Y. Miao, X. Weng, M. Wu, F. Yang, L.Zhou, F. Zhao, E. Chen: *Kineograph: Taking the Pulse of a Fast-Changing and Connected World*. EUROSYS 2012.

Aapo Kyrola, J. Bradley, D.Bickson and C. Guestrin. *Parallel Coordinate Descent for L1-Regularized Loss Minimization* (ICML 2011).

Y. Low, Aapo Kyrola, J. Gonzalez, D. Bickson, and C. Guestrin. *Machine Learning in the Cloud with GraphLab* (NIPS Workshop 2010 + preprint).

Y. Low, J. Gonzalez, Aapo Kyrola, D. Bickson, and C. Guestrin. *GraphLab: A New Framework for Parallel Machine Learning*. Presented in the conference for Uncertainty in Artificial Intelligence, Catalina Island, 2010.

## EDUCATION

**Carnegie Mellon University**, Pittsburgh, PA, USA.

**Ph.D., Computer Science** (2009 - 2014)

Advisors: Prof. Carlos Guestrin, Prof. Guy Blelloch.

Research areas:

external memory algorithms; large-scale machine learning, graph computation, recommendation systems and parallel optimization algorithms; distributed and parallel computing.

Graduate-level Courses:

Machine Learning (A), Computer Architecture (A-), Optimization (A-), Statistical Machine Learning (A), Algorithms in the Real World (A), Advanced Operating Systems (A), Semantics of Programming Languages (B), Graduate Algorithms(A-).

Teaching Assistant:

CS-210: Sequential and Parallel Data Structures and Algorithms, (Fall 2011, lectured by Guy Blelloch)

15-853: Algorithms in the Real World (Spring 2013, lectured by Guy Blelloch and myself)

**University of Helsinki**, Finland

**Bachelor of Social Sciences, Economics, September 2007**

including extended studies in Theoretical Physics and Mathematics

*Note: I conducted my studies while working part-time at my own company.*

## INTERNSHIPS

- Internship at Microsoft Research Asia, Systems Group. Mentored by Feng Zhao and Lidong Zhou. Worked on large-scale almost-realtime social network mining, resulting in an Eurosys 2012 paper. Summer 2011.
- Internship at the Personalization and Recommendations team, Twitter Inc, Fall 2012. Mentored by Pankaj Gupta. I worked mainly on algorithms to suggest people to follow based on their interests. During this project, I got good experience on the analytical challenges a “data scientist” faces and learned to use the Hadoop stack (particularly Pig and Scalding).

## SOFTWARE

- GraphChi, a disk-based large-scale graph computation engine: <http://graphchi.org>
- GraphLab, distributed graph computation system: <http://graphlab.org>

- Shotgun, parallel coordinate descent algorithm for solving L1-regularized problems: <http://www.select.cs.cmu.edu/projects/shotgun/>

## WORK EXPERIENCE

- Software Engineer, Facebook 2014 -
- Founder and Chief Owl, ProtoGeo (“Moves” app), 2012 - 2014
- Chief Scientist and Board Member, Sulake Corporation, 2001 - 2011
- Managing Director (co-founder), Sulake Corporation 2000 - 2001
- Senior Programmer, Satama Interactive (web-agency), 1998 - 2000
- Programmer (multimedia software), To the Point Oy, 1994 - 1998

### *Moves app / ProtoGeo*

I founded ProtoGeo with Sampo Karjalainen in May 2011 to develop a “gamified” exercise tracker for iPhone based on my idea. I developed the iPhone application and the backend for the application, which we launched for limited test use in October 2011. We quickly found out that the concept itself was not sufficiently appealing, but during the development of this app I learned the location technology stack of iOS in-depth, and based on Sampo’s idea, we started developing a fully automatic exercise tracker. Based on my design for an adaptive “hybrid” tracker, we were able to track user’s movements accurately enough for whole day, requiring user to charge only every night. After validating the basic technology, we decided to pivot and hire a founding team to develop the idea of “Moves”. My role in the new company was to first develop the first prototypes of the complete technology stack: iPhone tracker, server backend and the machine learning backend. After the early phase I have participated in various projects developing Moves towards its launch in January 2013, particularly in developing the activity recognition pipeline and algorithms. Our extremely successful launch led to huge increase in usage and urgent scalability challenges. I was instrumental in refactoring our backend to sustain high loads and designing a more scalable data storage model. I also ported the machine learning pipeline into C++, which allows it to be run directly in the phone (instead of the original version which processes the data in the server). Moves was selected by both Apple and Google as their top apps for year 2013. My title is entropy maximizing “Chief Owl”.

### *Habbo Hotel / Sulake*

I am the co-founder of **Sulake Corporation Inc.**, which is best known for its very successful **Habbo Hotel** (<http://www.habbo.com>) virtual world and community. Since the founding of the company in May 2000 (when I was 21 years old), over 200 million users have registered to our service. At its peak, Habbo attracted 15 million unique visitors each month and Sulake had offices

in 18 countries in five continents employing over 300 professionals. Sulake was selected as one of the fastest growing technology companies in EMEA region by Deloitte in 2006.

The original server and client softwares of Habbo were written by myself. I had also a strong influence into our business model and product planning. I was also the first managing director of the company before hiring a professional management. Before starting my Ph.D. studies, I was a leading developer in our new product team. I was member of the company board 2000-2011. My passion has always been to write great software, solve challenging technological problems in novel ways and find the best practices to facilitate efficient and creative development. I have always enjoyed working in a multidisciplinary team consisting of technology, creative and business professionals. I sold my shares to Elisa (Finnish telecommunication operator) in 2013.

#### *Early career*

I started working as a computer programmer already as a 15-year old, in a pioneering multimedia company To the Point Oy in Helsinki, Finland. The company produced mostly large multimedia CD-ROMs, which were used as marketing material or sold directly to customers. We used Macromedia Director software and its Lingo scripting language to create these projects, and I was quickly the leading programmer in the company, although I was simultaneously studying in high school. Although the projects were not technically very advanced, the huge size of them made the projects complicated. To enable designers to do their work efficiently, I had to write many kinds of tools and frameworks to facilitate the building of the projects. Here I learned that a good guiding principle for a programmer is to "learn hate manual work". Instead, try to automate everything and eliminate repetitive work, as it is prone to recurring mistakes.

After four years in To the Point, a team of us leaped to a prominent new web-agency Satama. In Satama I started doing mostly Java programming for work and learned about internet and database technologies. Many of the skills I learned in Satama enabled us to found Sulake in year 2000.

#### SPECIAL EXPERTISE

- **massive graph computation**
- **external memory algorithms**
- **large-scale machine learning and data analysis, parallel algorithms, convex optimization**
- **massively multiuser online services, distributed architectures, concurrent systems**
- message-driven architectures, object-relational database mapping and data access patterns

- API design, object-oriented design patterns
- agile software development processes

## AWARDS

- **VMWare Graduate Student Fellowship (2013)**
- **Nokia Foundation Recognition Award (2008)** for contributions in collaborative media and social networking
- **Finnish Engineering Award (2006)** for the development of the *Habbo Hotel* internet community, the Finnish Association of Graduate Engineers.
- **Finland prize (2005)** for the *Habbo Hotel* internet community, the Ministry of Education, Finland

## PROGRAMMING LANGUAGES & SYSTEMS

- Java (very fluent, since 1996)
- Scala (very good)
- C, C++, Objective-C (very good)
- Python (very good)
- OpenStreetMap map data
- Standard ML, Matlab
- Hadoop (MapReduce), MPI
- CUDA (GPU programming)
- other scripting languages: Lingo, ActionScript, JavaScript, Ruby
- others: HTML, CSS, Unix shell scripting

## LANGUAGES

- Finnish (mother tongue)
- English (fluent)
- Swedish (good)
- Chinese (moderate speaking skills, learning)
- German (workable)

## HOBBIES

- Go / "Weiqi" board game (ancient Chinese chess-like board game), rank: 2 kyu
- hiking in nature and bicycle trekking
- reading: literature, popular science, politics, current affairs
- Chinese language and culture

## MILITARY SERVICE

- Finnish armed forces, 6 months, motorized infantry, Karelian Brigade, Fall 2007.