Convergence Rate of Distributed Averaging Dynamics and Optimization in Networks

Angelia Nedić
University of Illinois at Urbana-Champaign
angelia@illinois.edu

the essence of knowledge
Boston — Delft

Foundations and Trends® in Systems and Control

Published, sold and distributed by:
now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:
now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

Angelia Nedić. *Convergence Rate of Distributed Averaging Dynamics and Optimization in Networks*. Foundations and Trends® in Systems and Control, vol. 2, no. 1, pp. 1–100, 2015.

This Foundations and Trends® issue was typeset in LaTeX using a class file designed by Neal Parikh. Printed on acid-free paper.

ISBN: 978-1-68083-041-5

Foundations and Trends® in Systems and Control

Volume 2, Issue 1, 2015

Editorial Scope

Topics

Foundations and Trends® in Systems and Control publishes survey and tutorial articles in the following topics:

- Control of:
  - Hybrid and discrete event systems
  - Nonlinear systems
  - Network systems
  - Stochastic systems
  - Multi-agent systems
  - Distributed parameter systems
  - Delay systems
- Systems
  - Energy storage
  - Grid integration
  - Conversion technologies
  - Underpinning materials developments
- Filtering, estimation, and identification
- Optimal control
- Systems theory
- Control applications

Information for Librarians

Foundations and Trends® in Systems and Control, 2015, Volume 2, 4 issues. ISSN paper version 2325-6818. ISSN online version 2325-6826. Also available as a combined paper and online subscription.

Foundations and Trends® in Systems and Control
Vol. 2, No. 1 (2015) 1–100

DOI: 10.1561/2600000004

Convergence Rate of Distributed Averaging Dynamics and Optimization in Networks

Angelia Nedić
University of Illinois at Urbana-Champaign
angelia@illinois.edu

Contents

Abstract

Recent advances in wired and wireless technology lead to the emergence of large-scale networks such as Internet, wireless mobile ad-hoc networks, swarm robotics, smart-grid, and smart-sensor networks. The advances gave rise to new applications in networks including decentralized resource allocation in multi-agent systems, decentralized control of multi-agent systems, collaborative decision making, decentralized learning and estimation, and decentralized in-network signal processing. The advances also gave birth to new large cyber-physical systems such as sensor and social networks. These network systems are typically spatially distributed over a large area and may consists of hundreds of agents in smart-sensor networks to millions of agents in social networks. As such, they do not possess a central coordinator or a central point for access to the complete system information. This lack of central entity makes the traditional (centralized) optimization and control techniques inapplicable, thus necessitating the development of new distributed computational models and algorithms to support efficient operations over such networks. This tutorial provides an overview of the convergence rate of distributed algorithms for coordination and its relevance to optimization in a system of autonomous agents embedded in a communication network, where each agent is aware of (and can communicate with) its local neighbors only. The focus is on distributed averaging dynamics for consensus problems and its role in consensus-based gradient methods for convex optimization problems, where the network objective function is separable across the constituent agents.

Angelia Nedić. *Convergence Rate of Distributed Averaging Dynamics and Optimization in Networks*. Foundations and Trends® in Systems and Control, vol. 2, no. 1, pp. 1–100, 2015.
DOI: 10.1561/2600000004.

1

Introduction

Recent technological advances resulted in various devices with computational and communication capabilities, such as a wide variety of sensors, robots, computers, laptops, cell-phones, iPads, etc. In turn, various networks of such interconnected devices have emerged giving birth to a wide range of physical and cyber systems. Some of these systems have been carefully engineered, while others have grown spontaneously on their own (such as Internet and many of the cyber-based social networks and data-base networks including Facebook, Twitter, Google, and YouTube). Due to the size of such networks, and often, due to the proprietary regulations, the complete network information is distributed among the entities that comprise the network and there is no central entity that controls or has access to the whole network information. In some networks, such as surveillance networks, centralized information architecture is often not desirable as it makes the system inoperable when the central entity fails. Instead, it is desirable to design the system with distributed information architecture in order to enhance the system robustness to a failure. Some of today's networks are mobile such as cell-phone networks or robotic networks. As such, they are characterized with a dynamic spatial-temporal connectivity structure. Thus, inherently, the algorithms for optimization and control of these networks have to be *distributed, robust, and adaptive to the time-varying connectivity structure* of the network.

Over the past decade, a substantial research effort has been spent in addressing the challenges imposed by such distributed and time-varying networks. Most relevant to this tutorial is the research on so called *consensus* (or agreement) problem studied within the control systems community. This problem deals with a quest of determining a decentralized control law that is compatible with the local agent knowledge of the network and that ensures the agent agreement on quantity (or a collection of quantities) asymptotically in time. The control laws solving the consensus problems, which are often referred to as *consensus protocols*, are at the core of the distributed algorithms that are discussed in this tutorial. In these algorithms, the consensus law is used as an underlying mechanism for diffusing the information from one agent to every other agent in the network.

Distributed computational models have a significant potential for affect several applications including distributed detection and estimation, machine learning, statistical inference, swarm intelligence, social networks, recommendation systems, computer systems, etc. Such computational models are relevant to any application where an aggregate behavior of a distributed networked agent system is to be monitored, estimated, or managed in order to achieve some system wide objective. These models are also useful where some aggregate system quantity is to be evaluated or estimated under restricted information access such as the absence of central entity, privacy-preserving restrictions, and partial and/or noisy observations of only a part of the network.

The goal of this tutorial to provide a connection between continuous optimization techniques from operations research and distributed averaging schemes to illustrate how standard optimization techniques could be used within multi-agent setting for distributed optimization over networks in a general setting. The main focus is on first-order optimization algorithms due to their low overhead computational cost and good stability properties with respect to noisy gradient evaluations, among others. In order to focus more on the interplay of the optimization techniques and distributed averaging schemes, the survey is focused on algorithms that do not use any knowledge about the global network properties such as, for example, the number of agents, the graph connectivity structure or any parameters characterizing connectivity (e.g., the second smallest eigenvalue of the associated graph Laplacian or a gap value of the associated Markov chain transition matrix). The algorithms that truly obey local connectivity structure are considered only, so the agents can only communicate with their local (one-hop) neighbors (i.e., multi-hop communications are not allowed). Furthermore, due to the inherent nature of iterative optimization algorithms, only a discrete-time setting is considered in this tutorial, while both static and time-varying graphs are addressed.

The basic discrete-time setting that we consider here corresponds to a so called *synchronous* update model, where all agents update at the same instances of time. There are many practical issues that arise for this model including difficulties associated with maintaining time-synchronization in face of computational and communication delays, reaching deadlock situations requiring re-initializations of the update process due to node/link failures. As our focus is on the interplay of two coupled processes (i.e., averaging dynamic and optimization), we are not going to spend significant time in addressing synchronicity issues. These issues can be resolved by using randomization techniques, as often done in communication networks by implementing a gossip- or broadcast-based updates (see, for example, papers by Boyd et al. [2005], Aysal et al. [2008], Nedić [2011b]), which are discussed to some extent in Chapter 5.

One of the main criticism of the distributed optimization algorithms that use weighted averaging protocols for information diffusion is that their convergence to an optimal solution of the problem of interest requires a construction of doubly stochastic matrices. This construction should be relying on local agent interactions and, moreover, it needs to be done at every step if the networks structure is time-varying. Simple and efficient such constructions exist for networks with bi-directional communication links even if the connectivity structure of the network is changing with time. However, such constructions are computationally prohibitive for networks with directed links, as shown by Gharesifard and Cortés [2012]. Some new algorithms have been recently developed that do not require doubly stochastic weights, as discussed in Chapter 5.

Another distributed algorithm for solving optimization problems in networks that has recently drawn significant attention is a distributed variant of the ADMM method (see survey by Boyd et al. [2010] for an elegant exposure of the ADMM). This algorithm is efficient for some structured problems since it is a dual-ascent algorithm. The major drawback of the current developments of distributed ADMM method is that it appears that the ADMM is limited to static bi-directional networks. It is a question whether the ADMM can be implemented in directed and/or time-varying networks. If it can be done, what is its efficiency? The author of this tutorial is unaware of any such work.

The tutorial is structured as follows. In Chapter 2 we provide an overview of consensus and distributed optimization problems. In Chapter 3, we focus on distributed weighted-averaging algorithms for solving the unconstrained consensus problem. Distributed algorithms for constrained consensus are discussed in Chapter 4. Basic distributed subgradient methods for solving optimization problems over networks is discussed in Chapter 5, while some conclusions are given in Chapter 6.

References

N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950.

R.J. Auman. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.

T.C. Aysal, M.E. Yildiz, A.D. Sarwate, and A. Scaglione. Broadcast gossip algorithms: Design and analysis for consensus. In *Proceedings of the 47th IEEE Conference on Decision and Control*, pages 4843–4848, 2008.

T.C. Aysal, M.E. Yildriz, A.D. Sarwate, and A. Scaglione. Broadcast gossip algorithms for consensus. *IEEE Transactions on Signal processing*, 57:2748–2761, 2009.

T. Başar, S.R. Etesami, and A. Olshevsky. Fast convergence of quantized consensus using Metropolis chains. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pages 1330–1334, 2014.

D. Bajović, J. Xavier, J.M.F. Moura, and B. Sinopoli. Consensus and products of random stochastic matrices: Exact rate for convergence in probability. *IEEE Transactions on Signal Processing*, 61(10):2557–2571, 2013.

H.H. Bauschke. *Projection Algorithms and Monotone Operators*. PhD thesis, Simon Frazer University, 1996.

H.H. Bauschke. Projection algorithms: Results and open problems. In D. Butnariu, Y. Censor, and S. Reich, editors, *Inherently Parallel Algorithms in Feasibility and Optimization and their Applications*, pages 11–22. Elsevier, Amsterdam, Netherlands, 2001.

H.H. Bauschke and J.M. Borwein. On projection algorithms for solving convex feasibility problems. *SIAM Review*, 38(3):367–426, 1996.

F. Benezit, V. Blondel, P. Thiran, J. Tsitsiklis, and M. Vetterli. Weighted gossip: distributed averaging using non-doubly stochastic matrices. In *Proceedings of the 2010 IEEE International Symposium on Information Theory*, pages 1753–1757, 2010.

D.P. Bertsekas. A note on error bounds for convex and nonconvex programs. *Computational Optimization and Applications*, 12:41–51, 1999.

D.P. Bertsekas and J.N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, Belmont, MA, 1997.

D.P. Bertsekas, A. Nedić, and A.E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, Belmont, Massachusetts, 2003.

D. Blackwell. Finite non-homogeneous chains. *Annals of Mathematics*, 46(4): 594–599, 1945.

P.A. Bliman and G. Ferrari-Trecate. Average consensus problems in networks of agents with delayed communications. *Automatica*, 44(8):1985–1995, 2008.

P.A. Bliman, A. Nedić, and A. Ozdaglar. Rate of convergence for consensus with delays. In *Proceedings of the 47th IEEE Conference on Decision and Control*, pages 4849–4854, 2008.

V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2996–3000, 2005.

S. Bolouki. *LINEAR CONSENSUS ALGORITHMS: STRUCTURAL PROPERTIES AND CONNECTIONS WITH MARKOV CHAINS*. PhD thesis, Univerité de Montréal, Département de Génie Électrique, 2014.

S. Bolouki and R.P. Malhamé. Theorems about ergodicity and class-ergodicity of chains with applications in known consensus models. In *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing*, pages 1425–1431, 2012.

V. Borkar and P.P. Varaiya. Asymptotic agreement in distributed estimation. *IEEE Transactions on Automatic Control*, 27(3):650–655, 1982.

V.S. Borkar, R. Makhijani, and R. Sundaresan. Asynchronous gossip for averaging and spectral ranking. *Selected Topics in Signal Processing*, 8(4): 703–716, 2014.

S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Gossip algorithms: Design, analysis, and applications. In *Proceedings of IEEE INFOCOM*, volume 3, pages 1653–1664, 2005.

S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.

S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010.

F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009.

M. Bürger, G. Notarstefano, F. Bullo, and F. Allgöwer. A distributed simplex algorithm for degenerate linear programs and multi-agent assignments. *Automatica*, 48(9):2298–2304, 2012.

J.V. Burke and M.C. Ferris. Weak sharp minima in mathematical programming. *SIAM Journal on Control and Optimization*, 31(6):1340–1359, 1993.

J.V. Burke and P. Tseng. A unified analysis of Hoffman's bound via Fenchel duality. *SIAM Journal on Optimization*, 6(2):265–282, 1996.

K. Cai and H. Ishii. Quantized consensus and averaging on gossip digraphs. *IEEE Transactions on Automatic Control*, 56(9):2087–2100, 2011.

M. Cao, D.A. Spielman, and A.S. Morse. A lower bound on convergence of a distributed network consensus algorithm. In *Proceedings of IEEE CDC*, pages 2356–2361, 2005.

M. Cao, A.S. Morse, and B.D.O. Anderson. Reaching a consensus in a dynamically changing environment: A graphical approach. *SIAM Journal on Control and Optimization*, 47(2):575–600, 2008a.

M. Cao, A.S. Morse, and B.D.O. Anderson. Reaching a consensus in a dynamically changing environment: Convergence rates, measurement delays, and asynchronous events. *SIAM Jounal on Control and Optimization*, 47 (2):601–623, 2008b.

M. Cao, A.S. Morse, and B.D.O. Anderson. Agreeing asynchronously. *IEEE Transactions on Automatic Control*, 53(8):1826–1838, 2008c.

R. Carli, F. Fagnani, A. Speranzon, and S. Zampieri. Communication constraints in coordinated consensus problem. In *Proceedings of IEEE American Control Conference*, pages 4189–4194, 2006.

R. Carli, F. Fagnani, P. Frasca, T. Taylor, and S. Zampieri. Average consensus on networks with transmission noise or quantization. In *Proceedings of European Control Conference*, pages 1852–1857, 2007.

R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. Gossip consensus algorithms via quantized communication. *Automatica*, 46(1):70–80, 2010.

F. Cattivelli and A.H. Sayed. Diffusion LMS strategies for distributed estimation. *IEEE Transactions on Signal Processing*, 58(3):1035–1048, 2010.

A. Cegielski and A. Suchocka. Relaxed alternating projection methods. *SIAM Journal on Optimization*, 19(3):1093–1106, 2008.

M. El Chamie and T. Başar. Optimal strategies for dynamic weight selection in consensus protocols in the presence of an adversary. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pages 735–740, 2014.

M. El Chamie, J. Liu, and T. Başar. Design and analysis of distributed averaging with quantized communication. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pages 3860–3865, 2014.

T-H. Chang, A. Nedić, and A. Scaglione. Distributed constrained optimization by consensus-based primal-dual perturbation method. *IEEE Transactions on Automatic Control*, 59(6):1524–1538, 2014.

B. Charron-Bost. Orientation and connectivity based criteria for asymptotic consensus. available at http://arxiv.org/abs/1303.2043, 2013.

J. Chen and A.H. Sayed. Diffusion adaptation strategies for distributed optimization and learning over networks. *IEEE Transactions on Signal Processing*, 60(8):4289–4305, 2012.

M.H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69:118–121, 1974.

F. Deutsch. Rate of convergence of the method of alternating projections. In B. Brosowski and F. Deutsch, editors, *Parametric Optimization and Approximation*, volume 76, pages 96–107. Birkhäuser, Basel, 1983.

F. Deutsch and H. Hundal. The rate of convergence for the cyclic projections algorithm II: Norms of nonlinear operators. *Journal of Approximation Theory*, 142:56–82, 2006a.

F. Deutsch and H. Hundal. The rate of convergence for the cyclic projections algorithm I: Angles between convex sets. *Journal of Approximation Theory*, 142:36–55, 2006b.

F. Deutsch and H. Hundal. The rate of convergence for the cyclic projections algorithm III: Regularity of convex sets. *Journal of Approximation Theory*, 155:155–184, 2008.

A.G. Dimakis, S. Kar, J.M.F. Moura, M.G. Rabbat, and A. Scaglione. Gossip algorithms for distributed signal processing. *Proceedings of the IEEE*, 98 (11):1847–1864, 2010.

A.D. Dominguez-Garcia and C. Hadjicostis. Distributed strategies for average consensus in directed graphs. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2124–2129, 2011.

J.C. Duchi, A. Agarwal, and M.J. Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3):592 –606, 2012.

S.R. Etesami and T. Başar. Convergence time for unbiased quantized consensus. In *Proceedings of the 52nd IEEE Conference on Decision and Control*, pages 6190–6195, 2013.

M.J. Fabian, R. Henrion, A.Y. Kruger, and J. Outrata. Error Bounds: Necessary and Sufficient Conditions. *Set-Valued Analysis*, 18:121–149, 2010.

F. Facchinei and J-S. Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*, volume I and II. Springer-Verlag New York, 2003.

F. Fagnani and S. Zampieri. Randomized consensus algorithms over large scale networks. *IEEE Journal on Selected Areas in Communications*, 26 (4):634–649, 2008.

M. Fiedler. *Special Matrices and Their Applications in Numerical Mathematics*. Dover Publications, Inc., New York, N. Y., USA, 2 edition, 2008.

B. Gharesifard and J. Cortés. Distributed strategies for generating weight-balanced and doubly stochastic digraphs. *European Journal of Control*, 18 (6):539–557, 2012.

B. Gharesifard and J. Cortés. Distributed continuous-time convex optimization on weight-balanced digraphs. *European Journal of Control*, 18(6): 539–557, 2012.

B. Gharesifard, T. Basar, and A. Dominguez-Garcia. Price-based distributed control for networked plug-in electric vehicles. In *Proceedings of the American Control Conference, Washington DC*, pages 5086–5091, 2013.

L.G. Gubin, B.T. Polyak, and E.V. Raik. The method of projections for finding the common point of convex sets. *USSR Computational Mathematics and Mathematical Physics*, 7(6):1–24, 1967.

I. Halperin. The product of projection operators. *Acta Scientiarum Mathematicarum*, 23:96–99, 1962.

Y. Hatano, A.K. Das, and M. Mesbahi. Agreement in presence of noise: Pseudogradients on random geometric networks. In *Proceedings of the 44th IEEE Conference on Decision and Control, and European Control Conference*, pages 6382–6387, 2005.

R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5, 2002.

J.M. Hendrickx. *Graphs and Networks for the Analysis of Autonomous Agent Systems*. PhD thesis, Université Catholique de Louvain, 2008.

J.M. Hendrickx and J.N. Tsitsiklis. A new condition for convergence in continuous-time consensus seeking systems. In *Proceedings of the 50th IEEE Conference on Decision and Control*, pages 5070–5075, 2011.

J.M. Hendrickx and J.N. Tsitsiklis. Convergence of type-symmetric and cut-balanced consensus seeking systems. *IEEE Transactions on Automatic Control*, 58(1):214–218, 2013.

A.J. Hoffman. On approximate solutions of systems of linear inequalities. *Journal of Research of the National Bureau of Standards*, 49(4):263–265, 1952.

M. Huang and J.H. Manton. Stochastic lyapunov analysis for consensus algorithms with noisy measurements. In *Proceedings of the 2007 IEEE American Control Conference*, pages 1419–1424, 2007a.

M. Huang and J.H. Manton. Stochastic approximation for consensus seeking: Mean square and almost sure convergence. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 306–311, 2007b.

M. Huang and J.H. Manton. Stochastic consensus seeking with measurement noise: Convergence and asymptotic normality. In *Proceedings of the 2008 IEEE American Control Conference*, pages 1337–1342, 2008.

M. Huang and J.H. Manton. Coordination and consensus of networked agents with noisy measurements: Stochastic algorithms and asymptotic behavior. *SIAM Journal on Control and Optimization*, 48(1):134–161, 2009.

L. Hurwicz. The design ofmechanisms for resource allocation. *American Economic Review*, 63:1–30, 1973.

A. Jadbabaie, J. Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.

D. Jakovetić, J. Xavier, and J.M.F. Moura. Cooperative convex optimization in networked systems: Augmented lagrangian algorithms with directed gossip communication. *IEEE Transactions on Signal Processing*, 59(8): 3889–3902, 2011.

D. Jakovetić, J. Xavier, and J.M.F. Moura. Fast distributed gradient methods. *IEEE Transactions on Automatic Control*, 59(5):1131–1146, 2014.

B. Johansson, M. Rabi, and M. Johansson. A simple peer-to-peer algorithm for distributed optimization in sensor networks. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 4705 –4710, Dec. 2007.

S. Kar and J.M.F. Moura. Distributed consensus algorithms in sensor networks: Link failures and channel noise. *IEEE Transactions on Signal Processing*, 57(1):355–369, 2009.

S. Kar and J.M.F. Moura. Distributed consensus algorithms in sensor networks: Quantized data and random link failures. *IEEE Transactions on Signal Processing*, 58(3):1383 –1400, 2010.

S. Kar and J.M.F. Moura. Convergence rate analysis of distributed gossip (linear parameter) estimation: Fundamental limits and tradeoffs. *IEEE Journal on Selected Topics on Signal Processing*, 5(4):674–690, 2011.

A. Kashyap, T. Basar, and R. Srikant. Quantized consensus. *Automatica*, 43 (7):1192–1203, 2007.

D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2003.

A. Kolmogoroff. Zur theorie der markoffschen ketten. *Mathematische Annalen*, 112(1):155–160, 1936.

A. Lalitha, T. Javidi, and A. Sarwate. Social learning and distributed hypothesis testing. *preprint arXiv:1410.4307*, 2015.

J. Lavaei and R.M. Murray. Quantized consensus by means of gossip algorithm. *IEEE Transactions on Automatic Control*, 57(1):19–32, 2012.

H. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram. Resilient asymptotic consensus in robust networks. *IEEE Journal on Selected Areas in Communications: Special Issue on In-Network Computation*, 31(4):766–781, 2013.

S. Lee and A. Nedić. Distributed random projection algorithm for convex optimization. *IEEE Journal on Selected Topics in Signal Processing*, 48 (6):988–1001, 2012.

S. Lee and A. Nedić. Asynchronous gossip-based random projection algorithms over networks. accepted in IEEE Transactions on Automatic Control, to appear in 2015, available at http://arxiv.org/pdf/1304.1757.pdf, 2013.

U. Lee and M. Mesbahi. Constrained consensus via logarithmic barrier functions. In *Proceedings of the 50th IEEE Conference on Decision and Control*, pages 3608–3613, 2011.

A. Lewis and J.-S. Pang. Error bounds for convex inequality systems. In J.-P. Crouzeix, J.-E. Martinez-Legaz, and M. Volle, editors, *Generalized Convexity, Generalized Monotonicity*, pages 75–110. Cambridge University Press, 1998.

N. Li and J. R. Marden. Designing games for distributed optimization. *IEEE Journal on Selected Topics in Signal Processing*, 7(2):230–242, 2013.

S. Li and T. Basar. Distributed learning algorithms for the computation of noncooperative equilibria. *Automatica*, 23(4):523–533, 1987.

J. Lin, A.S. Morse, and B.D.O. Anderson. The multi-agent rendezvous problem. In *Proceedings of the 42nd IEEE Conference on Decision and Control*, pages 1508–1513, 2003.

J. Lin, A.S. Morse, and B.D.O. Anderson. The multi-agent rendezvous problem - the asynchronous case. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, pages 1926–1931, 2004.

P. Lin and W. Ren. Distributed constrained consensus in the presence of unbalanced switching graphs and communication delays. In *Proceedings of the 51st IEEE Conference on Decision and Control*, pages 2238–2243, 2012.

Q. Ling and A. Ribeiro. Decentralized dynamic optimization through the alternating direction method of multiplier. *IEEE Transactions on Signal Processing*, 62(5):1185–1197, 2014.

J. Liu, A.S. Morse, B.D.O. Anderson, and C. Yu. Contractions for consensus processes. In *Proceedings of the 50th IEEE Conference on Decision and Control*, pages 1974–1979, 2011a.

J. Liu, A.S. Morse, B.D.O. Anderson, and C. Yu. Contractions for consensus processes. In *Proceedings of the 50th IEEE Conference on Decision and Control*, pages 1974–1979, 2011b.

J. Liu, S. Mou, A.S. Morse, B.D.O. Anderson, and C. Yu. Deterministic gossiping. *Proceedings of the IEEE*, 99(9):1505–1524, 2011c.

J. Liu, S. Mou, A.S. Morse, B.D.O. Anderson, and C. Yu. Deterministic gossiping. *Proceedings of the IEEE*, 99(9):1505–1524, 2011d.

J. Liu, A.S. Morse, A. Nedić, and T. Başar. Stability of a distributed algorithm for solving linear algebraic equations. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pages 3707–3712, 2014a.

J. Liu, A.S. Morse, A. Nedić, and T. Başar. Internal stability of linear consensus processes. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pages 922–927, 2014b.

J. Liu, A. Nedić, and T. Başar. Complex constrained consensus. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pages 1464–1469, 2014c.

I. Lobel and A. Ozdaglar. Distributed subgradient methods for convex optimization over random networks. *IEEE Transactions on Automatic Control*, 56(6):1291–1306, June 2011.

I. Lobel, A. Ozdaglar, and D. Feijer. Distributed multi-agent optimization with state-dependent communication. *Mathematical Programming*, 129(2): 255–284, 2011.

C. Lopes and A.H. Sayed. Distributed processing over adaptive networks. In *Adaptive Sensor Array Processing Workshop, MIT Lincoln Laboratory, MA*, pages 1–5, 2006.

C.G. Lopes and A.H. Sayed. Diffusion least-mean squares over adaptive networks: Formulation and performance analysis. *IEEE Transactions on Signal Processing*, 56(7):3122–3136, 2008.

J. Lorenz. A stabilization theorem for continuous opinion dynamics. *Physica A: Statistical Mechanics and its Applications*, 355:217–223, 2005.

J. Lorenz. *Repeated Averaging and Bounded Confidence-Modeling, Analysis and Simulation of Continuous Opinion Dynamics*. PhD thesis, Universität Bremen, 2007.

J. Lu and C.Y. Tang. Zero-gradient-sum algorithms for distributed convex optimization: the continuous-time case. *IEEE Transactions on Automatic Control*, 57(9):2348–2354, 2012.

Z.-Q. Luo. New error bounds and their applications to convergence analysis of iterative algorithms. *Mathematical Programming, Series B*, 88:341–355, 2000.

O.L. Mangasarian. Error bounds for nondifferentiable convex inequalities under strong Slater constraint qualification. Technical report, University of Wisconsin, 1996.

J.R. Marden, G. Arslan, and J.S. Shamma. Connections between cooperative control and potential games illustrated on the consensus problem. In *Proceedings of the European Control Conference*, pages 4604–4611, 2007.

A. Martinoli, F. Mondada, G. Mermoud, N. Correll, M. Egerstedt, A. Hsieh, L. Parker, and K. Stoy. *Distributed Autonomous Robotic Systems*. Springer Tracts in Advanced Robotics, Springer-Verlag, 2013.

A.S. Mathkar and V.S. Borkar. Nonlinear gossip. Preprint, 2014.

M. Mesbahi and M. Egerstedt. *Graph Theoretic Methods for Multiagent Networks.* Princeton University Press, Princeton, NJ, USA, 2010.

C.C. Moallemi and B. Van Roy. Consensus propagation. *IEEE Transactions on Information Theory*, 52(11):4753–4766, 2006.

L. Moreau. Stability of continuous-time distributed consensus algorithms. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, pages 3998–4003, 2004.

L. Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, 2005.

A. Nedić. Random projection algorithms for convex set intersection problems. In *Proceedings of the 49th IEEE Conference on Decision and Control*, pages 7655–7660, 2010.

A. Nedić. Random projection algorithms for convex minimization problems. *Mathematical Programming, Series B*, 129:225–253, 2011a.

A. Nedić. Asynchronous broadcast-based convex optimization over a network. *IEEE Transactions on Automatic Control*, 56(6):1337–1351, 2011b.

A. Nedić and J. Liu. Lyapunov approach to consensus problems. available at http://arxiv.org/abs/1407.7585, 2014.

A. Nedić and A. Olshevsky. Distributed optimization over time-varying directed graphs. available at http://arxiv.org/abs/1303.2289, 2013.

A. Nedić and A. Olshevsky. Stochastic gradient-push for strongly convex functions on time-varying directed graphs. available online at http://arxiv.org/abs/1406.2075, 2014.

A. Nedić and A. Ozdaglar. On the rate of convergence of distributed subgradient methods for multi-agent optimization. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 4711–4716, 2007.

A. Nedić and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009a.

A. Nedić and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009b.

A. Nedić and A. Ozdaglar. Cooperative distributed multi-agent optimization. In Y. Eldar and D. Palomar, editors, *Convex Optimization in Signal Processing and Communications*, pages 340–386. Cambridge University Press, 2010a.

A. Nedić and A. Ozdaglar. Convergence rate for consensus with delays. *Journal of Global Optimization*, 47(3):437–456, 2010b.

A. Nedić, A. Olshevsky, A. Ozdaglar, and J.N. Tsitsiklis. Distributed subgradient methods and quantization effects. In *Proceedings of the 47th IEEE Conference on Decision and Control*, pages 4177–4184, 2008.

A. Nedić, A. Olshevsky, A. Ozdaglar, and J.N. Tsitsiklis. On distributed averaging algorithms and quantization effects. *IEEE Transactions on Automatic Control*, 54(11):2506–2517, 2009.

A. Nedić, A. Ozdaglar, and P.A. Parrilo. Constrained consensus and optimization in multi-agent networks. *IEEE Transactions on Automatic Control*, 55(4):922–938, 2010.

A. Nedić, S. Lee, and M. Raginsky. Decentralized online optimization with global objectives and local communication. to appear in the forthcoming Proceedings of American Control Conference 2015, extended report available at http:http://maxim.ece.illinois.edu/pubs, 2014a.

A. Nedić, A. Olshevsky, and C.A. Uribe. Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs. to appear in the forthcoming Proceedings of American Control Conference 2015, extended report available at http://arxiv.org/abs/1410.1977, 2014b.

S. Oh, L. Schenato, P. Chen, and S. Sastry. Tracking and coordination of multiple agentsusing sensor networks: System Design, Algorithms and Experiments. *Proceedings of the IEEE*, 95(1):234–254, 2007.

R. Olfati-Saber and R.M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9):1520–1533, 2004.

R. Olfati-Saber, J.A. Fax, and R.M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

A. Olshevsky. *Efficient information aggregation for distributed control and signal processing*. PhD thesis, MIT, 2010.

A. Olshevsky and J.N. Tsitsiklis. Convergence rates in distributed consensus averaging. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 3387–3392, 2006.

A. Olshevsky and J.N. Tsitsiklis. On the nonexistence of quadratic lyapunov functions for consensus algorithms. *IEEE Transactions on Automatic Control*, 53(11):2642–2645, Dec. 2008.

A. Olshevsky and J.N. Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM Journal on Control and Optimization*, 48(1):33–55, 2009.

A. Olshevsky and J.N. Tsitsiklis. Degree fluctuations and the convergence time of consensus algorithms. *IEEE Transactions on Automatic Control*, 58(10):2626–2631, 2013.

S. Patterson and B. Bamieh. Distributed consensus with link failures as a structured stochastic uncertainty problem. In *Proceedings of the 46th Allerton Conference on Communication, Control, and Computing*, pages 623–627, 2008.

S. Patterson and B. Bamieh. Convergence rates of consensus algorithms in stochastic networks. In *Proceedings of the 49th IEEE Conference on Decision and Control*, pages 6608–6613, 2010.

S. Patterson, B. Bamieh, and A. El Abbadi. Distributed average consensus with stochastic communication failures. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 4215–4220, 2007.

S. Patterson, B. Bamieh, and A. El Abbadi. Distributed average consensus with stochastic communication failures. *IEEE Transactions on Signal processing*, 57:2748–2761, 2009.

B.T. Polyak. Random algorithms for solving convex inequalities. In D. Butnariu, Y. Censor, and S. Reich, editors, *Inherently Parallel Algorithms in Feasibility and Optimization and their Applications*, pages 409–422. Elsevier, Amsterdam, Netherlands, 2001.

S. Sundhar Ram, V.V. Veeravalli, and A. Nedić. Distributed and non-autonomous power control through distributed convex optimization. In *The 28th IEEE Conference on Computer Communications INFOCOM*, pages 3001–3005, 2009a.

S.S. Ram. *Distributed Optimization in Multi-agent Systems: Applications to Distributed Regression.* PhD thesis, University of Illinois at Urbana-Champaign, 2009.

S.S. Ram, A. Nedić, and V.V. Veeravalli. Distributed subgradient projection algorithm for convex optimization. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Taipei, Taiwan*, pages 3653–3656, 2009b.

S.S. Ram, A. Nedić, and V.V. Veeravalli. Asynchronous gossip algorithms for stochastic optimization. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC), Shanghai, China*, pages 3581–3586, 2009c.

S.S. Ram, A. Nedić, and V.V. Veeravalli. Distributed Stochastic Subgradient Projection Algorithms for Convex Optimization. *Journal of Optimization Theory and Applications*, 147:516–545, 2010a.

S.S. Ram, A. Nedić, and V.V. Veeravalli. Asynchronous gossip algorithms for stochastic optimization: Constant stepsize analysis. In *Recent Advances in Optimization and its Applications in Engineering*, volume the 14th Belgian-French-German Conference on Optimization (BFG), pages 51–60, 2010b.

S.S. Ram, A. Nedić, and V.V. Veeravalli. A new class of distributed optimization algorithms: application to regression of distributed data. *Optimization Methods and Software*, 27(1):71–88, 2012.

W. Ren. Consensus seeking in multi-vehicle systems with a time-varying reference state. In *Proceedings of IEEE American Control Conference*, pages 717–722, 2007.

W. Ren and R. Beard. Consensus seeking in multi-agent systems under dynamically changing interaction topologies. *IEEE Transactions on Automatic Control*, 50(5):655–661, 2005.

A.H. Sayed. Diffusion adaptation over networks. to appear in E-Reference Signal Processing, R. Chellapa and S. Theodoridis, editors, Elsevier, 2013. Also available online as arXiv:1205.4220v1, 2012.

A.H. Sayed. *Adaptation, Learning, and Optimization over Networks*, volume 7. Foundations and Trends in Machine Learning, 2014.

S. Shahrampour and A. Jadbabaie. Exponentially fast parameter estimation in networks using distributed dual averagingy. In *Proceedings of the 52nd IEEE Conference on Decision and Control*, pages 6196–6201, 2013.

S. Shahrampour, A. Rakhlin, and A. Jadbabaie. Distributed detection: Finite-time analysis and impact of network topology. *arXiv preprint arXiv:1409.8606*, 2014.

W. Shi, Q. Ling, G. Wu, and W. Yin. Extra: an exact first-order algorithm for decentralized consensus optimization. available online at:http://arxiv.org/abs/1404.6264, 2014.

K. Srivastava. *Distributed optimization with applications to sensor networks and machine learning*. PhD thesis, University of Illinois at Urbana-Champaign, Industrial and Enterprise Systems Engineering, 2011.

K. Srivastava and A. Nedić. Distributed asynchronous constrained stochastic optimization. *IEEE Journal of Selected Topics in Signal Processing*, 5(4): 772–790, 2011.

K. Srivastava, A. Nedić, and Dusan Stipanovic. Distributed constrained optimization over noisy networks. In *Proceedings of the 49th IEEE Conference on Decision and Control*, pages 1945–1950, 2010.

K. Srivastava, A. Nedić, and D. Stipanovic. Distributed bregman-distance algorithms for min-max optimization. In *Agent-Based Optimization*, pages 143–174. Springer Studies in Computational Intelligence (SCI), 2013.

S.S. Stanković, M.S. Stanković, and D.M. Stipanović. Consensus based overlapping decentralized estimator. *IEEE Transactions on Automatic Control*, 54(2):410–415, 2009.

S.S. Stanković, M.S. Stanković, and D.M. Stipanović. Decentralized parameter estimation by consensus based stochastic approximation. *IEEE Transactions on Automatic Control*, 56(3):531–543, 2011.

S. Sundaram and C.N. Hadjicostis. Distributed function calculation and consensus using linear iterative strategies. *IEEE Journal on Selected Areas in Communications: Issue on Control and Communications*, 26(4):650–660, 2008.

S. Sundaram and C.N. Hadjicostis. Distributed function calculation via linear iterative strategies in the presence of malicious agents. *IEEE Transactions on Automatic Control*, 56(7):1495–1508, 2011.

S. Sundaram, S. Revzen, and G.J. Pappas. A control-theoretic approach to disseminating values and overcoming malicious links in wireless networks. *Automatica*, 48(11):2894–2901, 2012.

A. Tahbaz-Salehi and A. Jadbabaie. A necessary and sufficient condition for consensus over random networks. *IEEE Transactions on Automatic Control*, 53(3):791–795, 2008.

A. Tahbaz-Salehi and A. Jadbabaie. Consensus over ergodic stationary graph processes. *IEEE Transactions on Automatic Control*, 55(1):225–230, 2010.

B. Touri. *Product of random stochastic matrices and distributed averaging.* PhD thesis, University of Illinois at Urbana-Champaign, Industrial and Enterprise Systems Engineering, 2011.

B. Touri. *Product of random stochastic matrices and distributed averaging.* Springer-Verlag, Berlin, 2012.

B. Touri and C. Langbort. On endogenous random consensus and averaging dynamics. *IEEE Transactions on Control of Network Systems*, 1(3):241–248, 2014.

B. Touri and A. Nedić. Distributed consensus over network with noisy links. In *Proceedings of the 12th International Conference on Information Fusion*, pages 146–154, 2009.

B. Touri and A. Nedić. When infinite flow is sufficient for ergodicity. *Proceedings of 49th IEEE Conference on Decision and Control*, pages 7479–7486, 2010.

B. Touri and A. Nedić. Approximation and limiting behavior of random models. *Proceedings of 49th IEEE Conference on Decision and Control*, pages 2656–2663, 2010.

B. Touri and A. Nedić. On ergodicity, infinite flow and consensus in random models. *IEEE Transactions on Automatic Control*, 56(7):1593–1605, 2011.

B. Touri and A. Nedić. On existence of a quadratic comparison function for random weighted averaging dynamics and its implications. In *Proceedings of the 50th IEEE Conference on Decision and Control*, pages 3806–3811, 2011a.

B. Touri and A. Nedić. Alternative characterization of ergodicity for doubly stochastic chains. In *Proceedings of the 50th IEEE Conference on Decision and Control*, pages 5371–5376, 2011b.

B. Touri and A. Nedić. On approximations and ergodicity classes in random chains. *IEEE Transactions on Automatic Control*, 57(11):2718–2730, 2012a.

B. Touri and A. Nedić. On backward product of stochastic matrices. *Automatica*, 48(8):1477–1488, 2012b.

B. Touri and A. Nedić. Product of random stochastic matrices and distributed averaging. *IEEE Transactions on Automatic Control*, 59(2):437–448, 2014.

P. Tseng. Successive projection under a quasi-cyclic order. Technical report, LIDS-P-1938, Massachusetts Institute of Technology, 1990.

K.I. Tsianos. *The role of the Network in Distributed Optimization Algorithms: Convergence Rates, Scalability, Communication / Computation Tradeoffs and Communication Delays*. PhD thesis, McGill University, Dept. of Electrical and Computer Engineering, 2013.

K.I. Tsianos and M.G. Rabbat. Distributed consensus and optimization under communication delays. In *Procedings of the 49th IEEE Allerton Conference on Communication, Control, and Computing*, pages 974–982, 2011.

K.I. Tsianos, S. Lawlor, and M.G. Rabbat. Consensus-based distributed optimization: Practical issues and applications in large-scale machine learning. In *Proceedings of the 50th Allerton Conference on Communication, Control, and Computing*, pages 1543–1550, 2012a.

K.I. Tsianos, S. Lawlor, and M.G. Rabbat. Push-sum distributed dual averaging for convex optimization. In *Proceedings of the 51st IEEE Conference on Decision and Control*, pages 5453–5458, 2012b.

J. N. Tsitsiklis and M. Athans. Convergence and asymptotic agreement in distributed decision problems. *IEEE Transactions on Automatic Control*, 29:42–50, 1984.

J.N. Tsitsiklis. *Problems in Decentralized Decision Making and Computation*. PhD thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1984.

J.N. Tsitsiklis, D.P. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.

S.-Y. Tu and A.H. Sayed. Diffusion strategies outperform consensus strategies for distributed estimation over adaptive networks. *IEEE Transactions on Signal Processing*, 60(12):6217–6234, 2012a.

S.-Y. Tu and A.H. Sayed. On the influence of informed agents on learning and adaptation over networks. *IEEE Transactions on Signal Processing*, 61(6):1339–1356, 2012b.

T. Vicsek, A. Czirok, E. Ben-Jacob, I. Cohen, and O. Schochet. Novel type of phase transitions in a system of self-driven particles. *Physical Review Letters*, 75(6):1226–1229, 1995.

J. von Neumann. *Functional Operators*. Princeton University Press, 1950.

P. Wan and M.D. Lemmon. Event-triggered distributed optimization in sensor networks. In *Symposium on Information Processing of Sensor Networks, (San Francisco, CA)*, pages 49–60, 2009.

J. Wang and N. Elia. A control perspective for centralized and distributed convex optimization. In *Proceedings of the IEEE Conference on Decision and Control, (Florida, USA)*, pages 3800–3805, 2011.

E. Wei and A. Ozdaglar. Distributed alternating direction method of multipliers. In *Proceedings of the 51st IEEE Conference on Decision and Control and European Control Conference*, pages 5445–5450, 2012.

E. Wei and A. Ozdaglar. On the O(1/k) convergence of asynchronous distributed alternating direction method of multipliers. In *Proceedings of IEEE Global Conference on Signal and Information Processing*, pages 551–554, 2013.

L. Xiao and S. Boyd. Fast linear iterations for distributed averaging. *Systems and Control Letters*, 53:65–78, 2004.

L. Xiao and S. Boyd. Optimal scaling of a gradient method for distributed resource allocation. *Journal of Optimization Theory and Applications*, 129 (3):469–488, 2006.

L. Xiao, S. Boyd, and S. Lall. A scheme for robust distributed sensor fusion based on average consensus. In *Proceedings of the 4th International Conference on Information Processing in Sensor Networks*, pages 63–70, 2005.

L. Xiao, S. Boyd, and S.-J. Kim. Distributed average consensus with least mean square deviation. *Journal of Parallel and Distributed Computing*, 67 (1):33–46, 2007.

F. Yan, S. Sundaram, S.V.N. Vishwanathan, and Y. Qi. Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties. *IEEE Transactions on Data and Knowledge Engineering*, 25(11): 2483–2493, 2013.

F. Zanella, D. Varagnolo, A. Cenedese, G. Pillonetto, and L. Schenato. Newton-raphson consensus for distributed convex optimization. In *Proceedings of the IEEE Conference on Decision and Control, (Florida, USA)*, pages 5917–5922, 2011.

B. Zhang, A. Lam, A. Dominguez-Garcia, and D. Tse. Optimal distributed voltage regulation in power distribution networks. to appear in IEEE Transaction on Power Systems, available at http://arxiv.org/abs/1204.5226, 2014.

M. Zhu and S. Martínez. On distributed convex optimization under inequality and equality constraints. *IEEE Transactions on Automatic Control*, 57(1): 151–164, 2012.

M. Zhu and S. Martínez. An approximate dual subgradient algorithm for distributed non-convex constrained optimization. *IEEE Transactions on Automatic Control*, 58(6):1534–1539, 2013.