Brief Announcement: Distributed Computing with Rules of Thumb

Aaron D. Jaggard*
Department of Computer Science, Colgate University
Hamilton, NY USA
DIMACS, Rutgers University
New Brunswick, NJ USA
adj@dimacs.rutgers.edu

Michael Schapira
Department of Computer Science
Princeton University
Princeton, NJ USA
ms7@cs.princeton.edu

Rebecca N. Wright†
Department of Computer Science and DIMACS
Rutgers University
New Brunswick, NJ USA
rebecca.wright@rutgers.edu

ABSTRACT

We present our recent work (ICS 2011) on dynamic environments in which computational nodes, or decision makers, follow simple and unsophisticated rules of behavior (*e.g.*, repeatedly "best replying" to others' actions, and minimizing "regret") that have been extensively studied in game theory and economics. We aim to understand when convergence of the resulting dynamics to an equilibrium point is guaranteed if nodes' interaction is not synchronized (*e.g.*, as in Internet protocols and large-scale markets). We take the first steps of this research agenda. We exhibit a general non-convergence result and consider its implications across a wide variety of interesting and timely applications: routing, congestion control, game theory, social networks and circuit design. We also consider the relationship between classical nontermination results in distributed computing theory and our result, explore the impact of scheduling on convergence, study the computational and communication complexity of asynchronous dynamics and present some basic observations regarding the effects of asynchrony on no-regret dynamics.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed applications*; F.1.1 [**Computation by Abstract Devices**]: Models of Computation

General Terms

Economics, Reliability, Theory

Keywords

Adaptive heuristics, game dynamics, convergence, self stabilization

1. MOTIVATION

Dynamic environments where computational nodes or human decision makers repeatedly interact arise in a variety of settings such as Internet protocols, large-scale markets, and multi-processor computer architectures. In many such settings, the behavior for the nodes is simple, natural and myopic—*i.e.*, they follow "rules of thumb" (or, in the language of Hart [5], "*adaptive heuristics*"). Often, this reflects the desire or necessity for nodes to provide quick responses and have a limited computational burden. Such behaviors, which include "repeated best response" and "regret minimization", have been extensively studied in game theory and economics.

In many interesting contexts, the resulting dynamics can, in the long run, move the global system in good directions and yield highly rational and sophisticated behavior, such as in game theory results demonstrating the convergence of best-response or no-regret dynamics to equilibrium points (see [5] and references therein). However, these positive results are based on the often unrealistic premise that nodes' actions are somehow synchronously coordinated.

It has long been known that asynchrony introduces substantial difficulties in distributed systems, as compared to synchrony [3], due to the "limitation imposed by local knowledge" [7]. There has been much work in distributed computing on identifying conditions that guarantee protocol termination in asynchronous computational environments. Over the past three decades, there have been many results regarding the possibility/impossibility borderline for failure-resilient computation (*e.g.*, [2, 7]). In these classical results, the risk of non-termination stems from the possibility of failures of nodes or other components. In contrast, in our setting the risk of non-convergence stems from limitations imposed by simplistic node behaviors.

We seek to bring together research in game theory and economics and research in distributed computing to form a new research agenda: "*distributed computing with adaptive heuristics*". Our aim is to investigate provable properties and possible worst-case system behavior of natural dynamics in asynchronous computational environments. We take the first steps of this research agenda in [6]. Here, we briefly describe the main contributions of [6]; we refer the reader to [6] for a more thorough exposition of our agenda, results, and directions for future research.

2. OUR CONTRIBUTIONS

We show that a large and natural class of dynamics fails to guarantee convergence to an equilibrium in an asynchronous

*Supported in part by NSF grants 0751674 and 0753492.
†Supported in part by NSF grant 0753061.

PODC'11, June 6–8, 2011, San Jose, California, USA.
ACM 978-1-4503-0719-2/11/06.

setting, even if the nodes and communication channels are reliably failure-free. This has implications across a wide domain of applications, ranging from routing and congestion control on the Internet to the stabilization of asynchronous circuits. Conversely, we show that non-convergence is not inherent to simple behaviors, as some forms of regret minimization provably converge in asynchronous settings. We also explore the impact of scheduling on convergence guarantees and the computational and communication complexity of asynchronous dynamics. In more detail, we make the following contributions:

Model.

We state a formal model for use in analyzing the convergence of myopic dynamics to equilibria in asynchronous computational environments. Our modeling approach draws on ideas both from work on dynamics in game theory and economics [5] and from work on self stabilization in distributed computing [1] and on protocol stability in networking [4].

General non-convergence result.

Due to practical constraints, it is often desirable or necessary that computational nodes' behavior rely on limited memory and processing power. In such contexts, nodes' decisions are often based on *bounded recall*—*i.e.*, dependent solely on recent history of interaction with others—and can even be *historyless*—*i.e.*, nodes only react to other nodes' current actions. We exhibit a general impossibility result showing that a broad class of bounded-recall behaviors, in which each node's behavior is additionally *self-independent*—*i.e.*, each node's choice of a new action is not allowed to depend on that node's previous actions—cannot always converge to a stable state. More specifically, we show that, for such behaviors, the existence of two "equilibrium points" implies that there is some execution that does not converge to any outcome even if all nodes and communication channels are guaranteed not to fail. We give evidence that our non-convergence result is essentially tight.

To prove our result we use a valency argument—a now-standard technique in distributed computing theory [2, 7], first introduced in the proof of the landmark non-termination result of Fischer, Lynch, and Paterson (FLP) for consensus protocols [3]. We point out that while the risk of protocol non-termination for consensus protocols in [3] and related work stems from the possibility of failures, the possibility of non-convergence in our framework stems from limitations on nodes' behaviors. Hence, there is no immediate translation from the FLP result to ours (and vice versa). We further discuss the link between consensus protocols and our framework in [6], where we also take an axiomatic approach and establish a non-termination result that holds for both contexts, thus unifying the treatment of these dynamic computational environments.

Implications across a wide variety of applications.

We apply our non-convergence result to a wide variety of interesting environments: (1) the Border Gateway Protocol (BGP), which handles Internet routing; (2) the Transmission Control Protocol (TCP), which handles congestion control on the Internet; (3) convergence of game dynamics to pure Nash equilibria; (4) stabilization of asynchronous circuits; and (5) diffusion of technologies in social networks.

Convergence, r-fairness, and randomness.

We study the effects on convergence to a stable state of natural restrictions on the order of nodes' activations (*i.e.*, the order in which nodes' have the opportunity to take steps) that have been extensively studied in distributed computing theory. We consider r*-fairness*, the guarantee that each node selects a new action at least once within every $r > 0$ consecutive time steps. Among other results, we show that there are systems that converge under all r-fair schedules, for some r that is exponential in the size of the system, but not for some $(r + 1)$-fair schedule.

Communication and computational complexity of asynchronous dynamics.

We study the tractability of determining whether convergence to a stable state is guaranteed (when such a guarantee is not precluded by our other results). We present two complementary hardness results establishing that, even for extremely restricted kinds of interactions, this task is hard: (1) an exponential communication complexity lower bound; and (2) a computational complexity PSPACE-completeness result that, alongside its computational implications, implies we cannot hope to have short witnesses of guaranteed asynchronous convergence (unless PSPACE = NP).

Asynchronous no-regret dynamics.

We present some basic observations about the convergence properties of no-regret dynamics in our framework, that establish that, in contrast to other simple behaviors, regret minimization is quite robust to asynchrony.

3. CONCLUSION

We believe that this work has but scratched the surface in the exploration of the behavior of natural dynamics in asynchronous computational environments. Many important questions remain wide open. In [6], we outline multiple general directions for future research.

4. REFERENCES

[1] S. Dolev. Self stabilization. *MIT Press*, 2000.

[2] F. Fich and E. Ruppert. Hundreds of impossibility results for distributed computing. *Distributed Computing*, 16(2–3):121–163, 2003.

[3] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.

[4] T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Trans. Netw.*, 10(2):232–243, 2002.

[5] S. Hart. Adaptive heuristics. *Econometrica*, 73:1401–1430, 2005.

[6] A. D. Jaggard, M. Schapira, and R. N. Wright. Distributed computing with adaptive heuristics. In *Proceedings of the Second Annual Symposium on Innovations in Computer Science*, pages 417–443, 2011.

[7] N. Lynch. A hundred impossibility proofs for distributed computing. In *Proceedings of the Eighth Annual ACM Symposium on Principles of Distributed Computing*, pages 1–28, 1989.