Implicit Coordination in Peer Production Networks[1]

JEFFREY V. NICKERSON, LEI ZHENG,
CHRISTIAN BOUWENSE AND FENG MAI Stevens Institute of Technology
DEBORAH M. GORDON Stanford University

Online peer production networks – networks that create artifacts like Wikipedia – are increasingly important for society. But they are strangely organized: They are notable for the absence of the explicit hierarchical command structures and functionally organized departments frequently seen in companies (Benkler, Shaw, & Hill, 2015). Researchers and practitioners alike are concerned about peer production's sustainability (Suh, Convertino, Chi, & Pirolli, 2009), and as a result want to understand it better (Butler, Joyce, & Pike, 2008; Geiger & Halfaker, 2013; Halfaker, Kittur, & Riedl, 2011; Panciera, Halfaker, & Terveen, 2009).

Recent studies of open source and peer production networks have noted that coordination in such networks can be explained at least in part by stigmergy (Bolici, Howison, & Crowston, 2016). The term stigmergy, coined by a researcher investigating the coordinated work of social insects (Grassé, 1960), refers to a process in which the traces of work become conditions or signals that generate more work. Social insect colonies operate without hierarchy or central control, yet groups of workers manage to build nests, collect food, raise brood, engage in defense, and to coordinate all of these activities in response to changing environmental conditions. Peer production is described as being stigmergic because coordination often happens not through explicit planning conversations, but through interactions triggered by previous interactions, all centered on the primary technical artifact, some form of text or source code (Bolici et al., 2016). Figure 1 shows traces of ant paths, on the left, and traces of peer production activity on the right (Han & Nickerson, 2015).

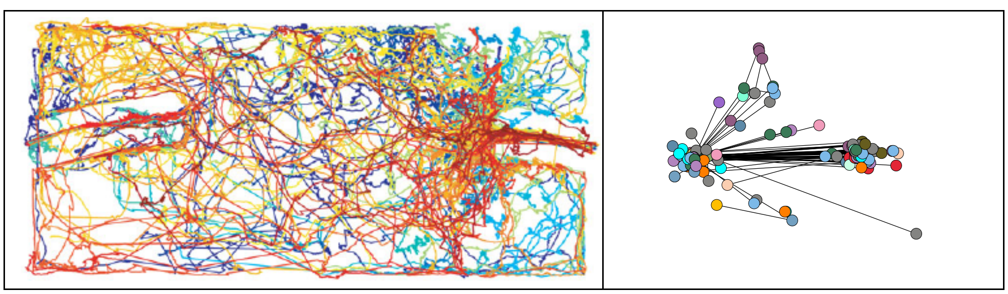

Figure 1. On the left, traces of ant paths (Pinter-Wollman, Wollman, Guetz, Holmes, & Gordon, 2011). On the right, traces of peer production.

Many studies have pointed out the potential application of stigmergy to human coordination (Blincoe & Damian, 2015; Cimino, Lazzeri, & Vaglini, 2015; Dorigo, Bonabeau, & Theraulaz, 2000; Elliott, 2016; Heylighen, 2016a, 2016b; Lewis & Marsh, 2016; Parunak, 2005; Robles, Merelo, & Gonzalez-Barahona, 2005; Susi, 2016). But to the best of our knowledge few if any of these past studies have involved biologists. This previous work might be extended and possibly transformed by applying the general theories and methods from mathematical biology that take the form of dynamic models (Ellner & Guckenheimer, 2006), as well as the specific dynamic models of the regulation of group activities by social insects (Gordon, 2002; Prabhakar, Dektar, & Gordon, 2012).

[1] This material is based upon work supported by the National Science Foundation under grants IIS-1422066, CCF-1442840, IIS-1717473, and IIS-1745463.

In the field of biology, some studies of stigmergy have focused on chemical communication in social insects. Social insects provide a system in which to investigate peer production. For example, ant colonies coordinate their work using interactions based on chemical cues. Ants smell with their antennae and in the course of a brief antennal contact, one ant can assess whether another is a nestmate, using the chemical profile of cuticular hydrocarbons on the other's body (Greene and Gordon 2003). Another form of chemical communication involves the detection of a chemical deposited on the substrate by another ant. This is a brief interaction with a time lag, since pheromones are volatile and disappear quickly. Each ant responds to the rate at which it meets other ants, or the amount of pheromone recently deposited by other ants, reflecting the rate at which other ants were at that location. In the aggregate, individual responses to the rate or recent pattern of such interactions produce the coordinated activity of colonies (Gordon, 2010).

Over more than 130 million years of evolution, ants have evolved diverse algorithms that regulate their work (Gordon, 2014, 2015). Such social algorithms may be analogous to processes at work, unrecognized, in human behavior. Once recognized, they might be improved. Furthermore, if some of these forms of coordination are not already at work in humans, they may suggest novel techniques for distributed coordination. In other words, social insects have over many years evolved sophisticated task allocation algorithms. Biologists have found ways of modeling the dynamics of these algorithms (Davidson et al. 2016; Gordon & Mehdiabadi, 1999), and both the models and the algorithms they depict may advance our understanding of peer production and lead to new techniques.

There is a need for new techniques. Studies of Wikipedia have shown that the growth rate of new articles has slowed (Suh et al., 2009). This is understandable; Wikipedia has been estimated to have involved more than 40,000,000 labor hours of effort as of 2013 (Geiger & Halfaker, 2013). But scientific knowledge appears to grow exponentially (He & Zhang, 2009), and so this slowdown in Wikipedia new article growth is concerning. Do we have examples of coordinated behavior in networks much larger than Wikipedia? In the animal world, some ant colonies have been estimated to contain more than 300,000,000 ants. Each large nest is the result of a number of labor hours orders of magnitude greater than our largest peer production project to date. While there may be limits to growth of an online encyclopedia, we don't know for sure if the slowdown in article growth is inevitable, or can be reversed through alternative forms of coordination. For example, simple shifts in coordination might reverse the current trend toward editing talk pages instead of articles. While many strong efforts have been made to increase skills and productivity on Wikipedia (Farzan & Kraut, 2013), and some patterns are emerging (Warncke-Wang, Ayukaev, Hecht, & Terveen, 2015), we still don't have a clear path to faster article growth. Moreover, since peer production is used on other large scale endeavors like the construction of software systems (Linux, Apache) and the design of physical objects (Thingiverse), designers of peer production systems would like to find ways of overcoming the current barriers to growth. In sum, the proposed work applies knowledge about coordination gained in biology to better understand, and to improve, coordination in peer production.

What are the temporal characteristics of Wikipedia edits? They exhibit burstiness (Goh and Barabasi 2008). This burstiness occurs at both the individual and group level; indeed, many articles are characterized by high editor concentration, meaning a few editors are responsible for most of the edits. Figure 2 shows a time series of edits and breaks it down by editor.

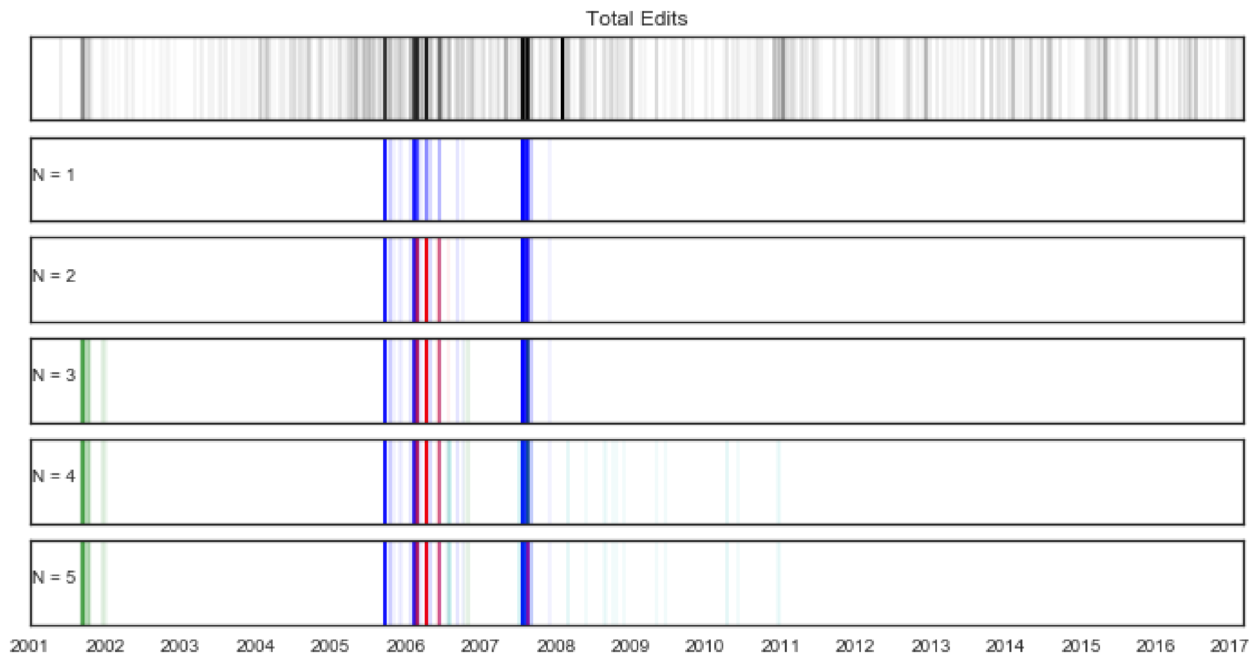

Figure 2. Edits of a featured article (NeXT) in Wikipedia between 2001 to 2017. The first row shows all edits, and the following rows decompose these into the edits by the five most prolific editors.

The analysis of this and other articles suggests that editing happens in bursts. Some of this burstiness may reflect the editing pattern of the most prolific editor, as in the second row of Figure 2. The editors do build on each other's work, as can be seen in the subsequent rows. Editor 2 builds on editor 1's work. Editor 3, the originator of the article, comes back to make changes after editors 1 and 2 have made substantial progress.

Burstiness may be a signal that attracts further coordination as well as a manifestation of past coordination. Just as frequency is used to coordinate ant behavior, it may also be used to coordinate human behavior. The techniques of mathematical biology may be applied to better understand such behavior.

Many studies of peer production have used models such as preferential attachment to explain the ways community members gravitate toward certain activities. Models such as preferential attachment, however, forecast unbounded growth. These models oversimplify the complexity of self-assignment of tasks. Models from mathematical biology may be useful, because in biology growth is often limited by competition and changes in the environment. Through evolution, animals have learned how to react to such changes. In particular, social insects have evolved algorithms for coordination. Because of this, dynamic models of task allocation of social insects might provide an alternative framework for studying coordination in peer production.

REFERENCES

Benkler, Y., Shaw, A., & Hill, B. M. (2015). Peer Production: A Modality of Collective Intelligence. In T. W. Malone & M. Bernstein (Eds.), The Handbook of Collective Intelligence. MIT Press.

Blincoe, K., & Damian, D. (2015). Implicit Coordination: A Case Study of the Rails Oss Project. IFIP International Conference on Open Source Systems.

Bolici, F., Howison, J., & Crowston, K. (2016). Stigmergic Coordination in Floss Development Teams: Integrating Explicit and Implicit Mechanisms. Cognitive Systems Research, 38, 14-22.

Butler, B., Joyce, E., & Pike, J. (2008). Don't Look Now, but We've Created a Bureaucracy: The Nature and Roles of Policies and Rules in Wikipedia. Proceedings of the SIGCHI conference on human factors in computing systems.

Cimino, M. G., Lazzeri, A., & Vaglini, G. (2015). Improving the Analysis of Context-Aware Information Via Marker-Based Stigmergy and Differential Evolution. International Conference on Artificial Intelligence and Soft Computing.

Davidson, J. D., Arauco-Aliaga, R. P., Crow, S., Gordon, D. M., & Goldman, M. S. (2016). Effect of Interactions between Harvester Ants on Forager Decisions. Frontiers in Ecology and Evolution, 4, 115.

Dorigo, M., Bonabeau, E., & Theraulaz, G. (2000). Ant Algorithms and Stigmergy. Future Generation Computer Systems, 16(8), 851-871.

Elliott, M. (2016). Stigmergic Collaboration: A Framework for Understanding and Designing Mass Collaboration. Mass Collaboration and Education (pp. 65-84): Springer.

Ellner, S., & Guckenheimer, J. (2006). Dynamical Models in Biology: Princeton University Press, Princeton, New Jersey, USA.

Farzan, R., & Kraut, R. E. (2013). Wikipedia Classroom Experiment: Bidirectional Benefits of Students' Engagement in Online Production Communities. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.

Geiger, R. S., & Halfaker, A. (2013). Using Edit Sessions to Measure Participation in Wikipedia. Proceedings of the 2013 conference on Computer supported cooperative work.

Goh, K.-I., & Barabási, A.-L. (2008). Burstiness and Memory in Complex Systems. EPL (Europhysics Letters), 81(4), 48002.

Gordon, D. M. (2002). The Organization of Work in Social Insect Colonies. Complexity, 8(1), 43-46.

Gordon, D. M. (2010). Ant Encounters: Interaction Networks and Colony Behavior. Princeton University Press.

Gordon, D. M. (2014). The Ecology of Collective Behavior. PLoS Biol, 12(3), e1001805.

Gordon, D. M. (2015). Collective Behavior in Animals: An Ecological Perspective. In T. W. Malone & M. Bernstein (Eds.), The Handbook of Collective Intelligence MIT Press.

Gordon, D. M., & Mehdiabadi, N. J. (1999). Encounter Rate and Task Allocation in Harvester Ants. Behavioral Ecology and Sociobiology, 45(5), 370-377.

Grassé, P. (1960). The Automatic Regulations of Collective Behavior of Social Insect and Stigmergy. Journal de psychologie normale et pathologique, 57(1), 1-10.

Greene, M. J., & Gordon, D. M. (2003). Social Insects: Cuticular Hydrocarbons Inform Task Decisions. Nature, 423(6935), 32-32.

Halfaker, A., Kittur, A., & Riedl, J. (2011). Don't Bite the Newbies: How Reverts Affect the Quantity and Quality of Wikipedia Work. Proceedings of the 7th international symposium on wikis and open collaboration.

Han, Y., & Nickerson, J. V. (2015). Exploring Design Space through Remixing. Collective Intelligence.

He, X., & Zhang, J. (2009). On the Growth of Scientific Knowledge: Yeast Biology as a Case Study. PLoS Comput Biol, 5(3), e1000320.

Heylighen, F. (2016a). Stigmergy as a Universal Coordination Mechanism I: Definition and Components. Cognitive Systems Research, 38, 4-13.

Heylighen, F. (2016b). Stigmergy as a universal coordination mechanism II: Varieties and evolution. Cognitive Systems Research, 38, 50-59.

Kittur, A., & Kraut, R. E. (2008). Harnessing the Wisdom of Crowds in Wikipedia: Quality Through Coordination. In Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (pp. 37–46). New York, NY, USA: ACM.

Lewis, T. G., & Marsh, L. (2016). Human Stigmergy: Theoretical Developments and New Applications. Cognitive Systems Research, 38, 1-3.

Panciera, K., Halfaker, A., & Terveen, L. (2009). Wikipedians Are Born, Not Made: A Study of Power Editors on Wikipedia. Proceedings of the ACM 2009 international conference on Supporting group work.

Parunak, H. V. D. (2005). A Survey of Environments and Mechanisms for Human-Human Stigmergy. International Workshop on Environments for Multi-Agent Systems.

Prabhakar, B., Dektar, K. N., & Gordon, D. M. (2012). The Regulation of Ant Colony Foraging Activity without Spatial Information. PLoS Comput Biol, 8(8), e1002670.

Robles, G., Merelo, J. J., & Gonzalez-Barahona, J. M. (2005). Self-Organized Development in Libre Software: A Model Based on the Stigmergy Concept. ProSim'05, 16.

Suh, B., Convertino, G., Chi, E. H., & Pirolli, P. (2009). The Singularity Is Not Near: Slowing Growth of Wikipedia. Proceedings of the 5th International Symposium on Wikis and Open Collaboration.

Susi, T. (2016). Social Cognition, Artefacts, and Stigmergy Revisited: Concepts of Coordination. Cognitive Systems Research, 38, 41-49.

Warncke-Wang, M., Ayukaev, V. R., Hecht, B., & Terveen, L. G. (2015). The Success and Failure of Quality Improvement Projects in Peer Production Communities. Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing.