



Emergent Collectives

Charles Petrie • Stanford University, retired

In recent columns, I've referred to an article I published in 2005 in a French journal on emergent collectives.¹ Though I've given the URL before (<http://www-cdr.stanford.edu/~petrie/revue/>), some new developments (including a new research project from Belgium [www.emergent-collectives.be/en/]), mean it's time to summarize and review that article's content here, as well as expand on it based on some recent experience.

Emergent Collectives Explain Disruptions

The article first reviews how certain distributed Internet-based systems have been unanticipated by people who should have been experts, such as RIAA executives, Bill Gates, ATM communications experts, Robert Metcalfe, and those in charge of France's Minitel system. I made the point that despite what Dilbert cartoons would have us believe, these people aren't fools, and asked what has been happening (and continues to happen) with these disruptions that surprised technically savvy people?

I can't resist digressing here. France has a long history of making iconic, wrong-headed Internet initiatives. Most recently, President Nicolas Sarkozy has proposed a "more civilized Internet" (see <http://tinyurl.com/3u946yy>). As one blogger recently said, good luck with that French Intranet (or they could just bring back the Minitel.) Here's a suggestion for Scott Adams: put Dilbert on loan to the French government to implement Sarkozy's suggestions.

Back to the question: How is it that all these smart folks have been so wrong-headed? A partial answer: engineers and government officials are trained to think of systems with a central control, and they're dealing with systems where the control is largely in the protocol.

This is an important point with respect to prediction and why I used the term "emergent." When a system's behavior is governed more by its protocol rather than some form of central control, that behavior is hard to predict with today's tools: we can't predict emergent behavior very well. And if you're predisposed to think in terms of central control, then certainly you'll be consistently fooled.

But that's not all. The Internet itself, the WWW, music and video file sharing, Wikipedia, and Linux have lots and lots of people contributing huge amounts of content and building out tremendous networks of information and functionality, for free! This goes counter to most older notions of economics.

Although I hadn't read "The Cathedral and the Bazaar" (www.catb.org/~esr/writings/cathedral-bazaar/cathedral-bazaar/) when I wrote the *Revue* article, Eric Raymond captures well the notion that people want to contribute to something larger than themselves. Social networks' success suggests that we abstract from this motivation because many postings at best can be described as building community by self-advertising. But people have certainly used social networks for good purposes, and the social motivation is present in any case.

The features of an emergent collective are

- a network of information/function nodes that has minimal central control, and that's largely controlled by a protocol specification,
- in which it's easy for people to add nodes to the network,
- and where they have a social incentive to do so.

Such networks grow quickly, and their behavior is difficult to predict, especially if your

models depend on central control and “rational economics.”

A Failed Prediction

What can we predict from this characterization, if it's true? Well, one of my related predictions from 2001 has somewhat failed 10 years later. I predicted that 3G cellular wouldn't be the “wireless Internet” of the future and Wi-Fi would (<http://www-cdr.stanford.edu/~petrie/802.11-Stockholm-2001/>). I wasn't entirely wrong in this, but I was wrong in the “emergent collectives” article to predict the growth of the user-powered Wi-Fi network, and it's instructive to consider this case.

I underestimated the power of cellular providers to throw unimaginable amounts of cash at the problem, including persuading people to be fearful of unsecured Wi-Fi networks, even persuading Germany and Italy to pass very restrictive connectivity legislation.

In the US, the cell providers have been remarkably successful at preventing cities from providing free Wi-Fi to their residents. Only Mountain View, California, has succeeded, and only because of Google. Finally, even I have a 3G smartphone, because it's the cheapest way for me to get Wi-Fi for my laptop at home, where I write this. The cellular providers have been much more successful at suppressing this particular emergent collective with the security/fear tactic than have been the record companies with the tactic of suing teenagers.

So consumer-provided free and open 802.11 networks have collapsed, largely due to consumers' fears that someone would sit in their driveway listening to their signals – even though for years any hackers inclined to sit in their driveway could easily crack the security most people used. People are no longer sharing their Wi-Fi, sadly leaving broadband providers with more control than necessary.

I speculate that Wi-Fi-sharing was a weak emergent collective because many people were sharing not because they wanted to contribute to something larger but because the routers came shipped with no security as a default.

Free and open hotspots have grown as businesses have increasingly recognized the economic benefits of doing so, much as I predicted in 2001. But the emergent collective of Wi-Fi has largely collapsed; it's instructive to see that the Wi-Fi network we have today isn't an emergent collective, and that emergent collectives can die, especially if their social incentive is weak.

What can we learn from this case, other than that prediction, especially of the future, is difficult? First, that the social incentive must be strong, and we don't know how to measure this. Second, we can't ignore traditional economic forces. We don't know how to measure this, either. So prediction just got harder.

Predicting Emergent Collectives

Apple's “apps” aren't an emergent collective as defined initially. They did create a platform in which it was relatively easy for individuals to add nodes and value. A sort of protocol exists in terms of the API. But the incentive is economic. It's *like* an emergent collective, but perhaps easier to predict because of the economic incentive. Maybe we could say there are two types: social and economic emergent collectives. The former remain harder to predict than the latter, but mixes will be even more difficult to predict.

Another prediction that I've made in this space² is that, because of the creeping commodization of everything, Amazon's Mechanical Turk (www.mturk.com/mturk/welcome) will become more mechanized and that, in fact, more complicated jobs will be done by “flash companies”

of individuals who link up for short time periods, supported by new Internet technologies that let them find each other (possibly via social networks) and coordinate their work.

Such a prediction might be more wishful thinking than an actuality, yet it has a good chance of coming true if only because of economic forces. We're all becoming self-employed. We'd better get used to it and at least hope for better Internet tool support. I still suspect such new tools will emerge, because a market exists for them: the increasing need to outsource larger and more complex tasks. This is being done right now by boutique consulting companies with access to various experts. This, too, will become more of a commodity owing to economic pressure. So good reasons exist to believe in such a future – but how can we really evaluate the likelihood of emergent collectives, either social or economic, when we have so few tools?

A Major Research Opportunity

Here's something that needs to happen in our research world: we should develop a better understanding of emergent behavior based on combinations of protocols and social incentives. Suppose we could design emergent collectives so that useful behavior resulted?

Were we able to do so, we could certainly achieve a lot more. For all the discussion in the distributed agents community about emergent behavior's value over the past 20 years (at least), very little in the way of engineering methods exist for predicting emergent behavior, much less designing it, although there is some related analysis.³ Were some researcher to actually develop a general algorithm that would predict collective behavior based on an interaction protocol and external constraints, it would be an important

development worthy of an international award.

Possibly the most important potential application of such new science would be in computational cognition. We still have no computational model of human minds. Such a model would explain, for instance, the phenomenon of attention and how it shifts, or how it affects subjective time. We don't have this model, but we do strongly suspect that our minds are composed of distributed (probably stupid) agents, largely unknown to our conscious mind (whatever that is), that somehow settle among themselves what gets brought to our conscious mind's attention, very unlike our computer operating systems' central time-sharing model.

Very early work in this area exists,⁴ but it doesn't yet tell us how attention shifts much less predict distributed agents' emergent behavior. A conference devoted to this research area will take place in November (<http://cogsys.org/acs/2011/home/>), and we can hope for important results.

A Practical Approach: Enterprise Simulation

We don't have to wait for such advanced science and discoveries to use the notion of emergent collectives to have more near-term impact. I suggested in the emergent collectives article¹ that companies perform simulations of possible emergent collectives to predict disruptions. Simulation is what we know how to do today: just let the distributed agents system run and see what happens.

I now suggest a very narrow but important practical application for enterprises, based on a recent experience that led me to view networks of people much as an Internet-based technology.

A lot of money and time is currently going into the topic of "innovative companies." But for most

companies, this is an oxymoron: they're designed to run routine processes efficiently and not for the consumer's benefit, despite mottoes to the contrary and laughably wrongly named "customer relations systems."

Yet, as John Hagel and John Seely Brown point out, these same companies are filled with smart, creative folks who go largely unrecognized (<http://blogs.hbr.org/bigshift/2010/04/are-all-employees-knowledge-wo.html>). Such people are frequently employed in routine jobs and must use their real talents in hobbies and "skunk projects" while the company suffers from being unable to respond adequately to changing market conditions and technologies. These two authors also discuss the formation of "creation networks" (www.johnhagel.com/paper_pushpull.pdf), which have some relation to emergent collectives but which focus on practical tactics companies might try to encourage productivity creativity.

Now imagine that networks of like-minded creative people inside companies could emerge that were designed to solve recognized problems by changing the company's behavior. Such systems would be kinds of games in which people would be motivated to creatively and collectively solve such problems. But we don't know enough about how to predict such emergent behavior, so how can we design such systems?

Again, we don't have to wait for new science and engineering. We can experiment. This isn't something firms tend to do – either they're doing poorly, in which case they don't want to waste resources, or they are doing well, in which case they don't think they need to improve. But companies can and should always be engaged in such small-scale experiments.

Google is running a small experiment by letting some employees work for one day per week on whatever they think is important. This is

contrary to the usual enterprise theory of effectiveness: Google is "wasting" 20 percent of its productivity. Yet, just looking in from the outside, this seems to be working, at least with respect to innovation. Wouldn't it be cool if companies could predict such outcomes prior to trying them for real?

Turns out, they can. They could simulate new behavior protocols in focus groups, with employees playing the roles of likely types of people. Such role-playing capabilities currently exist among management consultants to evaluate leadership effectiveness (http://mz-x.com/files/Info_LeadershipSimulation_en.pdf). Such simulations' capabilities could easily be adapted to test new internal business processes designed to address problems and increase creativity. We don't have to have the science to completely predict the emergent behavior in order to try out intuitions about how to change things for the better.

For instance, suppose you have a problem with mid-level managers not taking the initiative but rather always passing decisions up to the next level, which is a counter-innovative behavior. The solution is to simulate an incentive program that rewards decision-making. Problem with decision transparency? Build that into the game. Tweak as necessary. Include some typical role behaviors likely to prove problematic and see what happens.

I was a part of such a role-playing game development recently, and the results were quite surprising and rewarding to all involved in the exercise. One key insight was that the behavior wasn't determined by central fiat but rather by the interaction protocol, including both reward and social incentives. I wish I could talk about this more, but it was a small private exercise concerning a real problem in a real company. However, it gave me the confidence

to suggest that others experiment with this approach: it can be a practical way to change an enterprise's behavior.

Other than being pre-occupied with the day-to-day urgencies of getting the product out the door, there's no reason why companies can't conduct such role-playing simulations today, and there are very good reasons to do so.

My points here are that emergent collectives are an important feature of today's economy, that important fundamental research remains to be done in distributed (agent) systems, that these principles extend beyond Internet technologies into almost any group of people who nevertheless are connected by some protocol that governs their behavior, and that more could be done with existing tools. In particular,


role-playing simulations are a practical approach to designing protocols that can achieve new desired behaviors inside enterprises.

I look forward to new results in this early decade of the 21st century by today's young researchers and practitioners. □

References

1. C. Petrie, "Emergent Collectives for Work and Play," *AGIR Revue Generale de Strategie*, Societe de Strategie, La societe de 'information, nos. 20-21, 2005, pp. 146-152.
2. C. Petrie, "Plenty of Room Outside the Firm," *IEEE Internet Computing*, vol. 14, no. 1, 2010, pp. 92-96.
3. C. Rouff et al., "Properties of a Formal Method for Prediction of Emergent Behaviors in Swarm-Based Systems," *Proc. 2nd Int'l Conf. Software Eng. and Formal Methods (SEFM 04)*, IEEE Press, 2004, pp. 24-33.
4. W. Bridewell and P. Langley, "A Computational Account of Everyday Abductive Inference," *Proc. 33rd Ann. Meeting of the Cognitive Science Soc.*, Wiley, 2011; www.isle.org/~langley/papers/abduction.cogsci11.pdf.

Charles Petrie retired from Stanford University as a senior research scientist with the CS Logic Group. He received his PhD in computer science from the University of Texas at Austin. Petrie was a founding member of the technical staff of the MCC AI Lab, founding editor in chief of *IEEE Internet Computing*, founding executive director of the Stanford Networking Research Center, and founding chair of the Semantic Web Services Challenge. Contact him at petrie@stanford.edu.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.