

## OPINION



# Massively collaborative mathematics

The 'Polymath Project' proved that many minds can work together to solve difficult mathematical problems. **Timothy Gowers** and **Michael Nielsen** reflect on the lessons learned for open-source science.

On 27 January 2009, one of us — Gowers — used his blog to announce an unusual experiment. The Polymath Project had a conventional scientific goal: to attack an unsolved problem in mathematics. But it also had the more ambitious goal of doing mathematical research in a new way. Inspired by open-source enterprises such as Linux and Wikipedia, it used blogs and a wiki to mediate a fully open collaboration. Anyone in the world could follow along and, if they wished, make a contribution. The blogs and wiki functioned as a collective short-term working memory, a conversational commons for the rapid-fire exchange and improvement of ideas.

The collaboration achieved far more than Gowers expected, and showcases what we think will be a powerful force in scientific discovery — the collaboration of many minds through the Internet.

The specific aim of the Polymath Project was to find an elementary proof of a special case of the density Hales–Jewett theorem (DHJ), which is a central result of combinatorics, the branch of mathematics that studies discrete structures (see 'Multidimensional noughts and crosses'). This theorem was already known to be true, but for mathematicians, proofs are more than guarantees of truth: they are valued for their explanatory power, and a new proof of a theorem can provide crucial insights. There were two reasons to want a new proof of the DHJ theorem. First, it is one of a cluster of important related results, and although almost all the others have multiple proofs, DHJ had just one — a long and complicated proof

that relied on heavy mathematical machinery. An elementary proof — one that starts from first principles instead of relying on advanced techniques — would require many new ideas. Second, DHJ implies another famous theorem, called Szemerédi's theorem, novel proofs of which have led to several breakthroughs over the past decade, so there is reason to expect that the same would happen with a new proof of the DHJ theorem.

The project began with Gowers posting a description of the problem, pointers to background materials and a preliminary list of rules for collaboration (see [go.nature.com/DrCmnc](http://go.nature.com/DrCmnc)). These rules helped to create a polite, respectful atmosphere, and encouraged people to share a single idea in each comment, even if the idea was not fully developed. This lowered the barrier to contribution and kept the conversation informal.

## Building momentum

When the collaborative discussion kicked off on 1 February, it started slowly: more than seven hours passed before Jozsef Solymosi, a mathematician at the University of British Columbia in Vancouver made the first comment. Fifteen minutes later a comment came in from Arizona-based high-school teacher Jason Dyer. Three minutes after that Terence Tao (winner of a Fields Medal, the highest honour in mathematics) at the University of California, Los Angeles, made a comment. Over the next 37 days, 27 people contributed

approximately 800 substantive comments, containing 170,000 words. No one was specifically invited to participate: anybody, from graduate student to professional mathematician, could provide input on any aspect. Nielsen set up the wiki to distil notable insights from the blog discussions. The project received commentary on at least 16 blogs, reached the front page of the Slashdot technology-news

aggregator, and spawned a closely related project on Tao's blog. Things went smoothly: neither Internet 'trolls' — persistent posters of malicious or purposefully distracting comments

— nor well-intentioned but unhelpful comments were significant problems, although spam was an occasional issue on the wiki. Gowers acted as a moderator, but this involved little more than correcting a few typos.

Progress came far faster than anyone expected. On 10 March, Gowers announced that he was confident that the Polymath participants had found an elementary proof of the special case of DHJ, but also that, very surprisingly (in the light of experience with similar problems), the argument could be straightforwardly generalized to prove the full theorem. A paper describing this proof is being written up, along with a second paper describing related results. Also during the project, Tim Austin, a graduate student at the University of California, Los Angeles, announced another new (but non-elementary) proof of DHJ that made crucial use of ideas from the Polymath Project.

**"Who would have guessed that the working record of a mathematical project would read like a thriller?"**

The working record of the Polymath Project is a remarkable resource for students of mathematics and for historians and philosophers of science. For the first time one can see on full display a complete account of how a serious mathematical result was discovered. It shows vividly how ideas grow, change, improve and are discarded, and how advances in understanding may come not in a single giant leap, but through the aggregation and refinement of many smaller insights. It shows the persistence required to solve a difficult problem, often in the face of considerable uncertainty, and how even the best mathematicians can make basic mistakes and pursue many failed ideas. There are ups, downs and real tension as the participants close in on a solution. Who would have guessed that the working record of a mathematical project would read like a thriller?

**Broader implications**

The Polymath Project differed from traditional large-team collaborations in other parts of science and industry. In such collaborations, work is usually divided up in a static, hierarchical way. In the Polymath Project, everything was out in the open, so anybody could potentially contribute to any aspect. This allowed ideas to be explored from many different perspectives and allowed unanticipated connections to be made.

The process raises questions about authorship: it is difficult to set a hard-and-fast bar for authorship without causing contention or discouraging participation. What credit should be given to contributors with just a single insightful contribution, or to a contributor who is prolific but not insightful? As a provisional solution, the project is signing papers with a group pseudonym, ‘DHJ Polymath’, and a link to the full working record. One advantage of Polymath-style collaborations is that because all contributions are out in the open, it is



transparent what any given person contributed. If it is necessary to assess the achievements of a Polymath contributor, then this may be done primarily through letters of recommendation, as is done already in particle physics, where papers can have hundreds of authors.

The project also raises questions about preservation. The main working record of the Polymath Project is spread across two blogs and a wiki, leaving it vulnerable should any of those sites disappear. In 2007, the US Library of Congress implemented a programme to

preserve blogs by people in the legal profession; a similar but broader programme is needed to preserve research blogs and wikis.

New projects now under way will help to explore how collaborative mathematics works best (see [go.nature.com/4ZfIdc](http://go.nature.com/4ZfIdc)). One question of particular interest is whether the process can be scaled up to involve more contributors. Although DHJ Polymath was large compared with most mathematical collaborations, it fell short of being the mass collaboration initially envisaged. Those involved agreed that scaling up much further would require changes to the process. A significant barrier to entry was the linear narrative style of the blog. This made it difficult for late entrants to identify problems to which their talents could be applied. There was also a natural fear that they might have missed an earlier discussion and that any contribution they made would be redundant. In open-source software development, this difficulty is addressed in part by using issue-tracking software to organize development around ‘issues’ — typically, bug reports or feature requests — giving late entrants a natural starting point, limiting the background material that must be mastered, and breaking the discussion down into modules. Similar ideas may be useful in future Polymath Projects.

**Towards open science**

The Polymath process could potentially be applied to even the biggest open problems, such as the million-dollar prize problems of the Clay Mathematics Institute in Cambridge, Massachusetts. Although the collaborative model might deter some people who hope to keep all the credit for themselves, others could see it as their best chance of being involved in the solution of a famous problem.

Outside mathematics, open-source approaches have only slowly been adopted by scientists. One area in which they are being used is synthetic biology. DNA for the design of living organisms is specified digitally and uploaded to an online repository such as the Massachusetts Institute of Technology Registry of Standard Biological Parts. Other groups may use those designs in their laboratories and, if they wish, contribute improved designs back to the registry. The registry contains more than 3,200 parts, deposited by more than 100 groups. Discoveries have led to many scientific papers, including a 2008 study showing that most parts are not primitive but rather build on simpler parts (J. Peccoud *et al.* *PLoS ONE* 3, e2671; 2008). Open-source biology and open-source mathematics thus both show how science can be done using a gradual aggregation of

**Multidimensional noughts and crosses**

To understand the density Hales-Jewett theorem (DHJ), imagine a multidimensional noughts-and-crosses (or tic-tac-toe) board, with  $k$  squares on a side (instead of the usual three), and in  $n$  dimensions rather than two. Any square in this board has  $n$  coordinates between 1 and  $k$ , so for instance if  $k=3$  and  $n=5$ , then a typical point might be (1,3,2,1,2). A line on

such a board has coordinates that either stay the same from one point to the next, or go upwards or downwards. For instance, the three points (1,2,3,1,3), (2,2,3,2,2) and (3,2,3,3,1), form a line. DHJ states that, for a very large number of dimensions, filling in even a tiny fraction of the board always forces a line to be filled in somewhere — there is no possible way of

avoiding such a line. More than this, there is no way to avoid a ‘combinatorial line’, in which the coordinates that vary have to vary in the same direction (rather than some going up and some going down), as in the line (1,2,3,1,1), (2,2,3,2,2) and (3,2,3,3,3). The initial aim of the polymath project was to tackle the first truly difficult case of DHJ, which is when  $k=3$ .



insights from people with diverse expertise.

Similar open-source techniques could be applied in fields such as theoretical physics and computer science, where the raw materials are informational and can be freely shared online. The application of open-source techniques to experimental work is more constrained, because control of experimental equipment is often difficult to share. But open sharing

of experimental data does at least allow open data analysis. The widespread adoption of such open-source techniques will require significant cultural changes in science, as well as the development of new online tools. We believe that this will lead to the widespread use of mass collaboration in many fields of science, and that mass collaboration will extend the limits of human problem-solving ability. ■

**Timothy Gowers** is in the Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WB, UK, and a Royal Society 2010 Anniversary Research Professor. **Michael Nielsen** is a Toronto-based writer and physicist working on a book about the future of science. e-mails: W.T.Gowers@dpmms.cam.ac.uk; mn@michaelnielsen.org

## Stitching science together

Google Wave is the kind of open-source online collaboration tool that should drive scientists to wire their research and publications into an interactive data web, says **Cameron Neylon**.

Science communication today remains firmly wedded to its print origins. We cling to the notion that ‘the real version’ exists on the page. Beyond ease of delivery, we take very little advantage of the potential of the World Wide Web to transform the way we store and transfer knowledge. We rarely take the opportunity to update material with new data, or to provide a record of how a document or data set has changed. Gene names and protein structures should be routinely linked to database entries through hyperlinks. The outputs of computational processes should be connected to their inputs, so analyses can be redone. If we can make these records accessible to humans and readable by machines, then whole new types of analysis will become possible, indeed standard.

Many of these things are possible today. But they are hard to achieve. Much effort has gone into solving parts of the problem, by big players such as Microsoft and Amazon as well as by smaller organizations. Electronic lab notebooks can help to capture the details of science, and databases can make it available to the user. Reference-management tools such as Delicious, semantic data stores and Wikipedia can help to wire up and monitor knowledge. But the tools are often difficult to use and don’t ‘talk’ to each other. There is no single framework that makes it easy to link all the steps of science. Scientists do their analysis and writing using different software, and prepare graphs and record data using different tools.

Very few companies worldwide have both the expertise and resources to take on the task of stitching this together. So it is with great interest that I have watched Google develop its product, Google Wave. The company describes Google Wave as “what e-mail would look like if it were invented today”. It blends elements of e-mail with instant messaging and online collaborative authoring. The big change is that the

‘document’ or ‘wave’ is shared between all the participants and updates flow in real time. You no longer need to worry about which version of a document you have e-mailed around. This is helpful for scientists, but not revolutionary. Where Wave offers a big step for science is in two other functionalities.

### Two steps forward

First, Wave introduces the idea of robots: automated agents that can be invited into a document. Robots could look through your paper checking for Protein Data Bank codes or gene names, for example, and putting in links to the databases. A robot might represent a lab instrument, adding data automatically to your laboratory record when they become available. You can easily add maps, video or three-dimensional graphics to your work using ‘gadgets’ or ‘applications’, familiar from services such as iGoogle and Facebook. Robots can interact with this information, making it

possible to have a dashboard in your inbox to monitor and control instruments in the lab.

The second step forward is using versions. Each wave maintains a record of every change. It could be possible to check each step from data collection to drawing a graph and its publication. This would allow a reader to step through an analysis to see where conclusions have come from, and would make detecting fraud — or honest mistakes — much easier.

Google has done a good thing in making the protocol and programming tools open source, enabling people to test and build. Perhaps 50 people, myself included, from experimental scientists to journal publishers, have been testing the prototype system for science applications since June, building robots that link chemical information, visualize data and format references. Since 30 September, a much bigger group has been testing. But real benefits will come only if the system is widely adopted. Perhaps a new generation of scientists will be required to exploit the power that working with these dynamic documents and tools offers.

Solving the current problems in science communication requires the intervention of strong companies such as Google. But it will take more than technical advances to provoke scientists into taking full advantage of the web. We need pressure, and perhaps compulsion, from journals and funders to raise publishing standards to the new level made possible by such tools. Google Wave may not be, indeed is probably not, the whole answer. But it points the way to tools that build records and reproducibility into every step. And that has to be good for science. ■

**Cameron Neylon** is senior scientist in biomolecular sciences at the Science and Technology Facilities Council Rutherford Appleton Laboratory, Didcot OX11 0QX, UK. e-mail: cameron.neylon@stfc.ac.uk




ILLUSTRATION BY M. HODSON