

Stumbling Around in the Dark: Lessons from Everyday Mathematics

Ursula Martin^(✉)

University of Oxford, Oxford, UK
Ursula.Martin@cs.ox.ac.uk

Abstract. The growing use of the internet for collaboration, and of numeric and symbolic software to perform calculations it is impossible to do by hand, not only augment the capabilities of mathematicians, but also afford new ways of observing what they do. In this essay we look at four case studies to see what we can learn about the everyday practice of mathematics: the *polymath* experiments for the collaborative production of mathematics, which tell us about mathematicians attitudes to working together in public; the *minipolymath* experiments in the same vein, from which we can examine in finer grained detail the kinds of activities that go on in developing a proof; the mathematical questions and answers in *math overflow*, which tell us about mathematical-research-in-the-small; and finally the role of computer algebra, in particular the GAP system, in the production of mathematics. We conclude with perspectives on the role of computational logic.

1 Introduction

The popular image of a mathematician is of a lone genius (probably young, male and addicted to coffee) having a brilliant idea that solves a very hard problem. This notion has been fuelled by books such as Hadamard’s ‘Psychology of invention in the mathematical field’, based on interviews with forty or so mathematicians, and dwelling on an almost mystical process of creativity. Journalistic presentations of famous mathematicians continue the theme - for example picking out Andrew Wiles’s remark on his proof of Fermat’s conjecture “And sometimes I realized that nothing that had ever been done before was any use at all. Then I just had to find something completely new; it’s a mystery where that comes from.”

However Wiles also points out that inspiration is by no means the whole story, stressing the sheer slog of research mathematics:

“I used to come up to my study, and start trying to find patterns. I tried doing calculations which explain some little piece of mathematics. I tried to fit it in with some previous broad conceptual understanding of some part of mathematics that would clarify the particular problem I was thinking about. Sometimes that would involve going and looking it up in a book to see how it’s done there. Sometimes it was a question of modifying things a bit, doing a little extra calculation”

and draws attention to the lengthy periods of hard work between the moments of clarity:

“Perhaps I can best describe my experience of doing mathematics in terms of a journey through a dark unexplored mansion. You enter the first room of the mansion and it’s completely dark. You stumble around bumping into the furniture, but gradually you learn where each piece of furniture is. Finally, after six months or so, you find the light switch, you turn it on, and suddenly it’s all illuminated. You can see exactly where you were. Then you move into the next room and spend another six months in the dark. So each of these breakthroughs, while sometimes they’re momentary, sometimes over a period of a day or two, they are the culmination of – and couldn’t exist without – the many months of stumbling around in the dark that precede them.”

Cedric Villani also unpacks the myth in his recent book, which gives an account (incorporating emails, Manga comics, and his love of French cheese), of the work that won him the Field’s medal, and gives a gripping picture the mathematician’s days trying out ideas that don’t quite work, or turn out to be wrong, or right but not useful, and of the exhilaration of “the miracle” when “everything seemed to fit together as if by magic”. For Wiles, the process was essentially a solitary one, but for Villani “One of the greatest misconceptions about mathematics is that it’s a solitary activity in which you work with your pen, alone, in a room. But in fact, it’s a very social activity. You constantly seek inspiration in discussions and encounters and randomness and chance and so on.”

Wiles and Villani both talk about the importance of a broad view of mathematics, not so much for precise and formalisable correspondences, but in the hope that ideas that have worked in one area will stimulate new ways of looking at another. Edward Frenkel in “Love and Math”, his recent popular account his work on the Langlands Program, draws attention to the role of the past literature: “It often happens like this. One proves a theorem, others verify it, new advances in the field are made based on the new result, but the true understanding of its meaning might take years or decades”.

In a lecture in 2012 mathematician Michael Atiyah (who also remarked in the 1990’s that too much emphasis was placed on correctness, and mathematics needed a more “buccaneering” approach) pointed to the importance of errors in developing understanding: “I make mistakes all the time... I published a theorem in topology. I didn’t know why the proof worked, I didn’t understand why the theorem was true. This worried me. Years later we generalised it — we looked at not just finite groups, but Lie groups. By the time we’d built up a framework, the theorem was obvious. The original theorem was a special case of this. We got a beautiful theorem and proof.”

The personal accounts of outstanding mathematicians are complemented by the insights of ethnographers into the workaday worlds of less exalted individuals. Barany and Mackenzie, using the methods and language of sociology to look more keenly at this process of “stumbling around in the dark”, observe that

“the formal rigor at the heart of mathematical order becomes indissociable from the ‘chalk in hand’ character of routine mathematical work. We call attention to the vast labor of decoding, translating, and transmaterializing official texts without which advanced mathematics could not proceed. More than that, we suggest that these putatively passive substrates of mathematical knowledge and practice instead embody potent resources and constraints that combine to shape mathematical research in innumerable ways.”

Barany and Mackenzie observed mathematicians in their offices and in front of blackboards, but the growing use of the internet for collaboration, and numeric and symbolic software to perform calculations it is impossible to do by hand, not only augment the capabilities of mathematicians, in particular by enabling collaboration, but also provide new ways of observing what they do.

In this paper we look at four case studies to see what we can learn about the everyday practice of mathematics, so as to shed new light the process of “stumbling around in the dark”.

Two of the case studies are rooted in the mathematical area of group theory. A group is, roughly speaking, the set of symmetries of an object, and the field emerged in the nineteenth century, through the systematic study of roots of equations triggered by the work of Galois. It continues to provide a surprising and challenging abstract domain which underlies other parts of mathematics, such as number theory and topology, with practical applications in areas such as cryptography and physics. Its greatest intellectual achievement is the classification of finite simple groups, the basic “building blocks” of all finite groups. Daniel Gorenstein, one of the prime movers in coordinating the effort, estimates the proof variously as occupying between 5,000 and 15,000 journal pages over 30 years. Sociologist Alma Steingart describes the endeavour as “the largest and most unwieldy mathematical collaboration in recent history”, and points to the flexible notion of the idea of proof over the life of the collaboration (which explains the varied estimates of the length of the proof). She points out that the sheer volume of material meant that only one or two individuals were believed to have the knowledge to understand and check the proof, or to understand it well enough to fix the ‘local errors’ that were still believed to be present. The field has a well-developed tradition of computer support and online resources, in particular early heroic endeavours which constructed so called “sporadic” simple groups by constructing certain matrices over finite fields. Today widely used software such as GAP incorporates many specialist algorithms, and exhaustive online data resources, such as the ATLAS of data about simple groups, capture information about these complex objects.

Our four studies involve the *polymath* experiments for the collaborative production of mathematics, which tell us about mathematicians attitudes to working together in public; the *minipolymath* experiments in the same vein, from which we can examine in finer grained detail the kinds of activities that go on in producing a proof; the mathematical questions and answers in *mathoverflow*, which tell us about mathematical-research -in-the-small; and finally the role of computer algebra, in particular the GAP system, in the production of mathematics. We conclude with remarks on the role of computational logic.

2 The Power of Collaboration: *polymath*

Timothy Gowers was awarded a Fields Medal in 1998 for work combining functional analysis and combinatorics, in particular his proof of Szemerdi's theorem. Gowers has characterised himself as a problem-solver rather than a theory-builder, drawing attention to the importance of problem solvers and problem solving in understanding and developing broad connections and analogies between topics not yet amenable to precise unifying theories. He writes articulately on his blog about many topics connected with mathematics, education and open science, and used this forum to launch his experiments in online collaborative proof which he called "*polymath*". In a blog post on 27th January 2009 he asked "Is massively collaborative mathematics possible", suggesting that "If a large group of mathematicians could connect their brains efficiently, they could perhaps solve problems very efficiently as well.". Ground rules were formulated, designed to encourage massively collaborative mathematics both in the sense of involving as many people as possible: "we welcome all visitors, regardless of mathematical level, to contribute to active polymath projects by commenting on the threads"; and having a high degree of interaction and rapid exchange of informal ideas: "It's OK for a mathematical thought to be tentative, incomplete, or even incorrect". and "An ideal polymath research comment should represent a 'quantum of progress'."

The post attracted 203 comments from around the globe, exploring philosophical and practical aspects of working together on a blog to solve problems, and a few days later Gowers launched the first experiment. The problem chosen was to find a new proof of the density version of the "Hales Jewett Theorem", replacing the previously known very technical proof with a more accessible combinatorial argument which, it was hoped, would also open the door to generalisations of the result. Over the next seven weeks, 27 people contributed around 800 comments - around 170,000 words in all - with the contributors ranging from high-school teacher Jason Dyer to Gowers's fellow Fields Medallist Terry Tao. On March 10, 2009 Gowers was able to announce a new combinatorial proof of the result, writing "If this were a conventional way of producing mathematics, then it would be premature to make such an announcement - one would wait until the proof was completely written up with every single i dotted and every t crossed - but this is blog maths and we're free to make up conventions as we go along."

The result was written up as a conventional journal paper, with the author given as "D H J Polymath" - identifying the actual contributors requires some detective work on the blog - and published on the arxiv in 2009, and in the *Annals of Mathematics* in 2012. The journal version explains the process "Before we start working towards the proof of the theorem, we would like briefly to mention that it was proved in a rather unusual "open source" way, which is why it is being published under a pseudonym. The work was carried out by several researchers, who wrote their thoughts, as they had them, in the form of blog comments. Anybody who wanted to could participate, and at all stages of the process the comments were fully open to anybody who was interested".

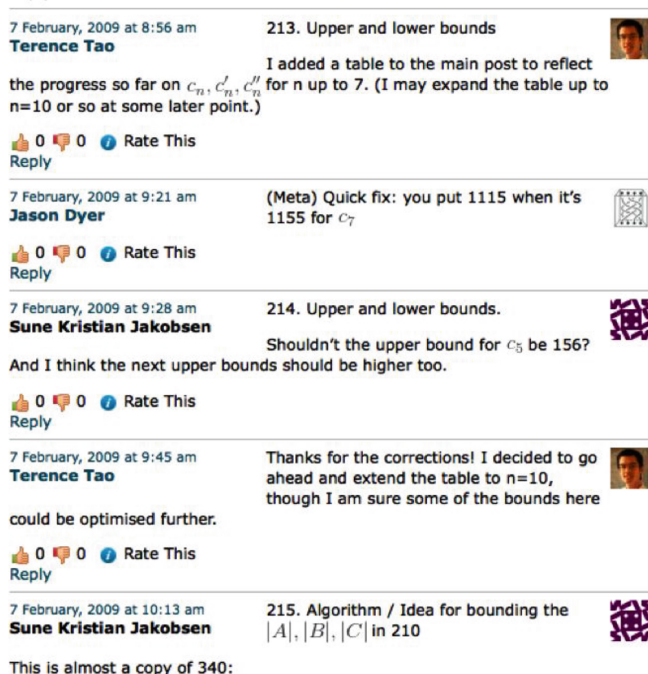


Fig. 1. An extract from the *polymath* blog for the proof of the Density Hales Jewett theorem

A typical extract from the blog (Fig. 1) shows the style of interaction. Participants, in line with the ground rules, were encouraged to present their ideas in an accessible way, to put forward partial ideas that might be wrong – “better to have had five stupid ideas than no ideas at all” – to test out ideas on other participants before doing substantial work on them, and to treat other participants with respect. As the volume of comments and ideas grew, it became apparent that the blog structure made it hard for readers to extract the thread of the argument and keep up with what was going on, without having to digest everything that had been previously posted, and in future experiments a leader took on the task of drawing together the threads from time to time, identifying the most appropriate next direction, and restarting the discussion with a substantial new blog post.

By 2015 there had been nine endeavours in the *polymath* sequence, and a number of others in similar style. Not all had achieved publishable results, with some petering out through lack of participation, but all have left the record of their partial achievements online for others to see and learn from - a marked contrast to partial proof attempts that would normally end up in a waste-basket.

The most recent experiment, *polymath* 8, was motivated by Yitang Zhang’s proof of a result about bounded gaps between primes. The twin primes conjecture states that there are infinitely many pairs of primes that differ by 2: 3, 5, . . . 11, 13

and so on. Zhang proved that there is a number K such that infinitely many pairs of primes differ by at most K , and showed that K is less than 70,000,000. After various discussions on other blogs, Tao formally launched the project, to improve the bound on K , on 13th June 2013. The first part of the project, *polymath 8a*, concluded with a bound of 4,680, and a research paper, also put together collaboratively, appeared on the arxiv in February 2014. The second part, *polymath 8b*, combined this with techniques independently developed in parallel by James Maynard, to reach a bound of 246, with a research paper appearing on the arxiv in July 2014. The participants also used Tao's blog to seek input for a retrospective paper reflecting on the experience, which appeared in the arxiv in September 2014.

One immediate concern was the scoping of the enquiry so as to not to intimidate or hamper individuals working on their own on this hot topic: it was felt that this was more than countered by providing a resource for the mathematical community that would capture progress, and provide a way to pull together what would otherwise be many independent tweaks. The work was well suited to the *polymath* approach: the combination of Tao's leadership and the timeliness of the problem made it easy to recruit participants; the bound provided an obvious metric of progress and maintained momentum; and it naturally fell into five components forming what Tao called a "factory production line". The collaborative approach allowed people to learn new material, and brought rapid sharing of expertise, in particular knowledge of the literature, and access to computational and software skills.

Tao himself explained how he was drawn into the project by the ease of making simple improvements to Zhang's bound, even though it interrupted another big project (a piece of work that he expects to take some years), and summed up by saying "All in all, it was an exhausting and unpredictable experience, but also a highly thrilling and rewarding one." The time commitment was indeed intense - for example, a typical thread "*polymath 8b*, II: Optimising the variational problem and the sieve" started on 22 November 2013, and ran for just over two weeks until 8th December. The initial post by Tao runs to about 4000 words - it is followed by 129 posts, of which 36, or just under a third, are also by Tao.

Tao and other participants were motivated above all by the kudos of solving a high-profile problem, in a way that was unlikely had they worked individually, but also by the excitement of the project, and the enthusiasm of the participants and the wider community. They reported enjoying working in this new way, especially the opportunity to work alongside research stars, the friendliness of the other participants, and their tolerance of errors, and the way in which the problem itself and the *polymath* format provided the incentive of frequent incremental progress, in a way not typical of solo working.

Participants needed to balance the incentives for participation against other concerns. Chief among these was the time commitment: participants reported the need for intense concentration and focus, with some working on it at a "furious pace" for several months; some feeling that the time required to grasp everything that was happening on the blog make *polymath* collaborations more,

rather than less, time consuming than traditional individual work or small-group collaboration; and some feeling that the fast pace was deterring participants whose working style was slower and more reflective.

Pure mathematicians typically produce one or two journal papers a year, so that, particularly for those who do not yet have established positions, there will be concerns that a substantial investment of time in *polymath* might damage their publication record. While such a time commitment would normally be worth the risk for a high-profile problem that has the likely reward of a good publication, the benefits are less clear-cut when the paper is authored under a group pseudonym (D H J Polymath), with the list of participants given in a linked wiki. As a participant remarked “*polymath* was a risk for those who did not have tenure”. On the other hand, in a fast moving area, participants may feel that incorporating their ideas into the collective allows them to make a contribution that they would not have achieved with solo work, or that engaging in this way is better than being beaten to a result and getting no credit, especially if participation in a widely-read blog is already adding to their external reputation.

An additional risk for those worried about their reputation can be that mistakes are exposed for ever in a public forum: pre-tenure mathematician Pace Nielsen was surprised that people were “impressed with my bravery” and would advise considering this issue before taking part. Rising star James Maynard observed: “It was very unusual for me to work in such a large group and so publicly - one really needed to lose inhibitions and be willing to post ideas that were not fully formed (and potentially wrong!) online for everyone to see.”

Those reading the *polymath* 8 sites went well beyond the experts - with an audience appreciating the chance to see how mathematics was done behind the scenes, or as Tao put it “How the sausage is made”. At its height it was getting three thousand hits a day, and even readers who knew little mathematics reported the excitement of checking regularly and watching the bounds go down. All the members of a class of number theory students at a summer school on Bounded Gaps admitted to following *polymath*. Perhaps typical was Andrew Roberts, an undergraduate who thanked the organisers for such an educational resource and reported “reading the posts and following the ‘leader-board’ felt a lot like an academic spectator sport. It was surreal, a bit like watching a piece of history as it occurred. It made the mathematics feel much more alive and social, rather than just coming from a textbook. I don’t think us undergrads often get the chance to peek behind closed doors and watch professional mathematicians ‘in the wild’ like this, so from a career standpoint, it was illuminating.” David Roberts, an Australian educator who used *polymath* in his classes to show students how the things they were learning were being used in cutting-edge research, reported “For me personally it felt like being able to sneak into the garage and watch a high-performance engine being built up from scratch; something I could never do, but could appreciate the end result, and admire the process.” The good manners remarked upon by the expert participants extended to less-well informed users, with questions and comments from non-experts generally getting a polite response, often from Tao himself, and a few more outlandish comments, such as claims of a simple proof, being ignored except for a plethora of down-votes.

The experiments have attracted widespread attention, in academia and beyond. Gowers had worked closely with physicist Michael Nielsen in designing the wiki and blog structure to support *polymath*, and in an article in *Nature* in 2009 the pair reflected on its wider implications, a theme developed further in Nielsen's 2012 book, *Reinventing Discovery*, and picked up by researchers in social and computer science analysing the broader phenomenon of open science enabled by the internet.

Like the participants, the analysts remarked on the value of the *polymath* blogs for capturing the records of how mathematics is done, the kinds of thinking that goes into the production of a proof, such as experimenting with examples, computations and concepts, and showing the dead ends and blind alleys. As Gowers and Nielsen put it, "Who would have guessed that the working record of a mathematical project would read like a thriller?"

Although *polymath* is often described as "crowdsourced science", the crowd is a remarkably small and expert one. The analogy has often been drawn with open source software projects - however these are typically organised in a much more modular and top down fashion than is possible in developing a mathematical proof, where many ideas and strands will be interwoven in a manner, as Nielsen comments, much more akin to a novel.

Research on collaboration and crowdsourcing carried out by psychologists, cognitive scientists and computer scientists helps explain the success of *polymath* and other attempts at open collaboration in mathematics. Mathematicians have well-established shared standards for exposition and argument, making it easy to resolve disputes. As the proof develops, the blog provides a shared cognitive space and short term working memory. The ground rules allow for a dynamic division of labour, and encourage a breakdown into smaller subtasks thus reducing barriers to entry and increasing diversity. At the same time, presenting the whole activity to readers, rather than in a more rigidly structured and compartmentalised way, allows more scope for serendipity and conversation across threads. Gowers gives an example of one contributor developing ideas in a domain which he was not familiar with (ergodic theory), and another who translated these ideas into one that he was familiar with (combinatorics), thus affecting his own line of reasoning.

A striking aspect of *polymath* is that senior figures in the field are prepared to try such a bold experiment, to think though clearly for themselves what the requirements are, and to take a "user centred" view of the design, based on their understanding of their own user community. For example it was suggested that participants might use a platform such as github, designed for collaborative working and version control, to make the final stage, collaborating on a paper, more straightforward. Tao responded "One thing I worry about is that if we use any form of technology more complicated than a blog comment box, we might lose some of the participants who might be turned off by the learning curve required."

These working records allow the analysis of what is involved in the creation of a proof, which we explore in the following section.

3 Examples, Conjectures, Concepts and Proofs: *minipolymath*

The *minipolymath* series applied the *polymath* model to problems drawn from the International Mathematical Olympiad (IMO), a competition for national teams of high school students. Tao and Gowers are among successful Olympiad contestants who have gone on to win Fields medals. Using short but challenging high-school level problems allowed for a much greater range of participants, and for greater experimentation with the format, as problems typically took hours rather than months to be solved. This windmill-inspired problem was composed by Geoff Smith for the 2011 IMO, held in the Netherlands.

The windmill Problem. Let S be a finite set of at least two points in the plane. Assume that no three points of S are collinear. A *windmill* is a process that starts with a line l going through a single point $P \in S$. The line rotates clockwise about the pivot P until the first time that the line meets some other point Q belonging to S . This point Q takes over as the new pivot, and the line now rotates clockwise about Q , until it next meets a point of S . This process continues indefinitely.

Show that we can choose a point P in S and a line l going through P such that the resulting windmill uses each point of S as a pivot infinitely many times.
(Tao, 8:00 pm)

Tao posted the problem as a *minipolymath* challenge at 8 pm on July 19th, 2011 a few days after the competition. Interest was immediate, and seventy four minutes later the participants had found a solution and by 9.50 pm, when Tao called a halt there were 147 comments on the blog, over 27 threads. To investigate this we developed a typology of comments as below (nine comments fell into two categories, and one fell into three). Figure 2 shows the proportion of each category.




Fig. 2. The proportions of comments which concerned conjectures, concepts, proofs, examples, and other.

Exemplars of the Typology of Comments

Concept. Since the points are in general position, you could define “the wheel of p ”, $w(p)$ to be radial sequence of all the other points $p \neq p$ around p . Then, every transition from a point p to q will “set the windmill in a particular spot” in q . This device tries to clarify that the new point in a windmill sequence depends (only) on the two previous points of the sequence. (Anonymous, 8:41 pm)

Example. If the points form a convex polygon, it is easy. (Anonymous, 8:08 pm)

Conjecture. One can start with any point (since every point of S should be pivot infinitely often), the direction of line that one starts with however matters! (Anonymous, 8:19 pm)

Conjecture. Perhaps even the line does not matter! Is it possible to prove that any point and any line will do? (Anonymous, 8:31 pm)

Proof. The first point and line P_0, l_0 cannot be chosen so that P_0 is on the boundary of the convex hull of S and l_0 picks out an adjacent point on the convex hull. Maybe the strategy should be to take out the convex hull of S from consideration; follow it up by induction on removing successive convex hulls. (Haggai Nuchi, 8:08 pm)

Example and Conjecture. Can someone give me *any* other example where the windmill cycles without visiting all the points? The only one I can come up with is: loop over the convex hull of S . (Srivatsan Narayanan, 9:08 pm)

Other. Got it! Kind of like a turn number in topology. Thanks! :) (Gal, 9:50 pm)

Analysing the typology in more depth we see that:

- *Examples* played a key role in the discussion, forming around a third of the comments. We saw some supporting examples to explain or justify conjectures, concepts and requests for clarification; but most of the examples concerned counter examples to the conjectures of others, or explanations to why these counter examples were not in fact, counter examples. As an IMO problem, the question was assumed to be correctly stated, but a number of ‘counterexample’ comments concerned participants attempts to understand it, with a number of people initially misled by the windmill analogy into thinking of the rotating line as a half-line, in which case the result does not hold, and a counterexample can indeed be found.
- *Conjectures.* Conjectures concerned possible translations to other domains which would provide results that could be applied; extensions of the initial problem; sub-conjectures towards a proof; and conjectures of properties of the windmill process aimed at understanding it better and clarifying thinking.
- *Concepts.* One class of concepts concerned analogies to everyday objects: as well as the somewhat misleading windmills, these included

“We could perhaps consider “layers” of convex hulls (polygons) .. like peeling off an onion. If our line doesn’t start at the “core” (innermost)

polygon then I feel it'll get stuck in the upper layers and never reach the core." (Varun, 8:27 pm)

Notice the author's use of apostrophes to stress that this is an analogy. Other analogies were to related mathematical objects, so as to provide ideas or inspiration, rather than affording an exact translation so that results from the new domain could be immediately applied. An important development was the emergence of the idea of talking about the "direction" of a line, leading to the observation, important for the final proof, that the number of points on each side of the line stays constant throughout the windmill process. It was noticeable how new concepts rapidly spread among the participants, even before they were precisely pinned down, enabling communication.

- *Proof*. Twenty one comments concerned a proof. Fourteen were about possible proof strategies, one was clarification of the strategy, one was carrying out a plan and five were about identifying which properties were relevant to the proof. Three strategies were discussed: one by induction and two involving translation to analogous domains. Within this the proof itself occupies a mere 16 comments, by 7 participants (three comments are "Anonymous" but from the context appear to be the same person).
- *Other*. The 34 comments classified as "other" include clarification of duplicate comments, explanations of a claim, and friendly interjections. Some of these are mathematically interesting, guiding the direction of the discussion, while others are simply courtesy comments. All play an important role, along with smiley faces, exclamation marks, and so on, in creating and maintaining an environment which is friendly, collaborative, informal and polite.

4 Questions and Answers: *mathoverflow*

Discussion fora for research mathematics have evolved from the early newsgroup newsgroups to modern systems based on the *stackexchange* architecture, which allow rapid informal interaction and problem solving. In three years *mathoverflow.net* has hosted 61,000 conversations and accumulated over 10,000 users, of whom about 500 are active in any month. The highly technical nature of research mathematics means that this is not currently an endeavour accessible to the public at large: a separate site *math.stackexchange.com* is a broader question and answer site "for people studying math at any level and professionals in related fields".

Within *mathoverflow* house rules give detailed guidance, and stress clarity, precision, and asking questions with a clear answer. Moderation is fairly tight, and some complain it constrains discussion. The design of such systems has been subject to considerable analysis by the designers and users, and *meta.mathoverflow* contains many reflective discussions. A key element of the success of the system is user ratings of questions and responses, which combine to form reputation ratings for users. These have been studied by psychologists

Tausczik and Pennebaker who concluded that *mathoverflow* reputations offline (assessed by numbers of papers published) and in *mathoverflow* were consistently and independently related to the *mathoverflow* ratings of authors' submissions, and that while more experienced contributors were more likely to be motivated by a desire to help others, all users were motivated by building their *mathoverflow* reputation.

We studied the mathematical content of *mathoverflow* questions and responses, choosing the subdomain of group theory so as to align with related work on GAP: at the time of writing (April 2015) around 3,500 of the *mathoverflow* questions are tagged "group theory", putting it in the top 5 topic-specific tags.

We analysed a sample of 100 questions drawn from April 2011 and July 2010 to obtain a spread and developed a typology:

Conjecture 36 % — asks if a mathematical statement is true. May ask directly "Is it true that" or ask under what circumstances a statement is true.

What is this 28 % — describes a mathematical object or phenomenon and asks what is known about it.

Example 14 % — asks for examples of a phenomenon or an object with particular properties.

Formula 5 % — ask for an explicit formula or computation technique.

Different Proof 5 % — asks if there is an alternative to a known proof. In particular, since our sample concerns the field of group theory, a number of questions concern whether a certain result can be proved without recourse to the classification of finite simple groups.

Reference 4 % — asks for a reference for something the questioner believes to be already in the literature.

Perplexed 3 % — ask for help in understanding a phenomenon or difficulty. A typical question in this area might concern why accounts from two different sources (for example Wikipedia and a published paper) seem to contradict each other.

Motivation 3 % — asks for motivation or background. A typical question might ask why something is true or interesting, or has been approached historically in a particular way.

Other 2 % — closed by moderators as out of scope, duplicates etc.

We also looked for broad phenomena in the structure of the successful responses. *mathoverflow* is very effective, with 90 % of our sample successful, in that they received responses that the questioner flagged as an "answer", of which 78 % were reasonable answers to the original question, and a further 12 % were partial or helpful responses that moved knowledge forward in some way. The high success rate suggests that, of the infinity of possible mathematical questions, questioners are becoming adept at choosing those for *mathoverflow* that are amenable to its approach. The questions and the answers build upon an assumption of a high level of shared background knowledge, perhaps at the level of a PhD in group theory.

The usual presentation of mathematics in research papers is in a standardised precise and rigorous style: for example, the response to a conjecture is either a counterexample, or a proof of a corresponding theorem, structured by means of intermediate definitions, theorems and proofs. By contrast, the typical response to a *mathoverflow* question, whatever the category, is a discussion presenting facts or short chains of inference that are relevant to the question, but may not answer it directly. The facts and inference steps are justified by reference to the literature, or to mathematical knowledge that the responder expects the other participants to have. Thus in modelling a *mathoverflow* discussion, we might think of each user as associated to a collection of facts and short inferences from them, with the outcome of the discussion being that combining the facts known to different users has allowed new inferences. Thus the power of *mathoverflow* comes from developing collective intelligence through sharing information and understanding.

In 56% of the responses we found citations to the literature. This includes both finding papers that questioners were unaware of, and extracting results that are not explicit in the paper, but are straightforward (at least to experts), consequences of the material it contains. For example, the observation needed from the paper may be a consequence of an intermediate result, or a property of an example which was presented in the paper for other purposes. In 34% of the responses, explicit examples of particular groups were given, as evidence for, or counter examples to, conjectures. The role of examples in mathematical practice, for example as evidence to refine conjectures, was explored by Lakatos: we return to this below.

In addition *mathoverflow* captures information known to individuals but not normally recorded in the research literature: for example unpublished material, motivation, explanations as to why particular approaches do not work or have been abandoned, and intuition about conjectures. The presentation is often speculative and informal, a style which would have no place in a research paper, reinforced by conversational devices that are accepting of error and invite challenge, such as “I may be wrong but...”, “This isn’t quite right, but roughly speaking...”. Where errors are spotted, either by the person who made them or by others, the style is to politely accept and correct them: corrected errors of this kind were found in 37% of our sample (we looked at *discussions* of error: we have no idea how many actual errors there are).

It is perhaps worth commenting on things that we did not see in our sample of technical questions tagged “group theory” in *mathoverflow*. In developing “new” mathematics considerable effort is put into the formation of new concepts and definitions: we saw little of this in *mathoverflow*, where questions are by and large focussed on extending or refining existing knowledge and theories. A preliminary scan suggests these are not present in other technical areas of *mathoverflow* either.

We see little serious disagreement in our *mathoverflow* sample: perhaps partly because of the effect of the “house rules”, but also because of the style of discussion, which is based on evidence from the shared research background and

knowledge of the participants: there is more debate in *meta.mathoverflow*, which has a broader range of non-technical questions about the development of the discipline and so on.

5 Everyday Calculation: GAP

GAP (Groups, Algorithms and Programming) is a substantial open-source computer algebra system, supporting research and teaching in computational group and semigroup theory, discrete mathematics, combinatorics and finite fields, and the applications of these techniques in areas such as cryptography and physics. It has been developed over the past 20 years or so by teams led initially from the University of Aachen, and currently from the University of St Andrews. According to google scholar it has been cited in around 3,500 research papers: the GAP making list has over a thousand members.

GAP provides well documented implementations of algorithms covering techniques for identifying, and computing in, finite and infinite groups defined by permutations, generators and relations, and matrices over finite fields. It also supports a variety of standard data-sets: for example the 52×10^{12} semigroups with up to 10 elements. It currently comprises over 0.6 million lines of code, with a further 1.1 million in over 100 contributed packages. Considerable effort is taken to ensure that GAP packages and datasets can be treated as objects in the scholarly ecosystem through establishing refereeing standards, citation criteria and so on.

Alongside the efforts one would expect of running a large open source project - a source code repository, mailing lists and archives, centralized testing services, issue tracker, release management, and a comprehensive website - the activity of the core GAP developers is driven by extending the power and reach of the system. Thus extensive efforts are being put into techniques for increasing the efficiency of algorithms handling matrices, permutations, finite fields and the like, for example by developing new data representations, and parallelising GAP over multicores in ways that do not increase complexity for the user. Extending the reach of GAP includes developing new theories and algorithms, and supporting these with high quality well-documented implementations: for example recent work has included devising computational methods for semigroups and new techniques for computing minimal polynomials for matrices over finite fields.

Research users of GAP typically use it to experiment with conjectures and theories. Whereas pencil and paper calculation restricts investigations to small and atypical groups, the ready availability in GAP of a plethora of examples, and the ease of computing with groups of large size, makes it possible to develop, explore and refine hypotheses, examples and possible counter examples, before proceeding to decide exactly what theorems to prove, and developing the proofs in a conventional journal paper. For example, we reviewed the 49 papers in google scholar which cited GAP Version 4.7. 5, 2014. A number of items were eliminated: duplicates; out of scope, such as lecture slides; and papers that did not appear to cite GAP, or cited it without mention in the text. The remaining 37 papers fell into six main groupings:

Explicit Computation as Part of a Proof, 25 % — These papers contained proofs that needed explicit and intricate calculation, carried out in GAP but difficult or impossible to do by hand. This arises particularly in theorems that depend on aspects of the classification of finite simple groups, or other results of a similar character, and hence require checking a statement for an explicitly given list of groups, each of which can be handled in GAP.

Examples and Counter examples, 25 % — These papers had used GAP to find or verify explicit examples of groups or other combinatorial objects: in some cases to illustrate a theorem, or as evidence for a conjecture; in others as counter-examples to a conjectured extension or variant of a result. Notice that GAP's built-in libraries of groups are often used to search for counter examples.

New Algorithms, 20 % — These were papers mainly devoted to the exposition of a new algorithm. In some cases these were supported by an explicit GAP implementation. In the rest the algorithm was more general than could be supported by GAP, but the paper contained a worked example, executed in GAP, for illustrative purposes.

Computations with Explicit Primes, 14 % — Groups whose number of elements is a power of a prime are the basic “building blocks” of finite group theory. As we have indicated, GAP can compute with fixed values of a prime number p , but is unable to handle statements of the form “For all primes p”. Many results of this form have generic proofs for “large enough” primes, while requiring a different proof for fixed small values of p , which can be computed by GAP, sometimes by making use of GAP's built in tables of particular families of groups. Thus for example Vaughan-Lee finds a formula for the number of groups of order p^8 , with exponent p , where p is prime, $p > 7$. To complement this he computes a list of all 1,396,077 groups of order 3^8 , and these are made available in GAP.

Applications in Other Fields, 10 % — This included three papers in theoretical physics and gauge theory, all doing explicit computation using GAPs built in representations of Lie Algebras, and an example of how GAP has made group theory accessible to non-specialists. A further paper in education research presented symmetry and the Rubik's cube.

Other, 6 % — Two papers cited GAP as background material, one in describing how their own algorithmic approach went beyond it; and one mentioning that calculations in GAP had shown a claimed result in an earlier paper to be incorrect, and presenting a corrected statement and proof.

One factor encouraging the take-up of GAP in research is its widespread use in teaching mathematics, both at undergraduate level and as part of research training. For example a professor at Colorado State University in the USA writes “I have been using GAP for many years in my undergraduate and graduate classes in algebra and combinatorics [...]. I have found the system an indispensable tool for illustrating phenomena that are beyond simple pencil-and-paper methods [...]. It also has been most useful as a laboratory environment for students to investigate algebraic structures [...]. [T]his first-hand investigation gives

students a much better understanding of what these algebraic structures are, and how their elements behave, than they would get by the traditional examples presented in a board-lecture situation.” GAP enables an experimental approach, where students can explore examples and formulate and solve research questions of their own, developing their skills as mathematicians and building familiarity and confidence in using tools such as GAP later in their careers.

The examples above highlight the use of GAP in published mathematics research: supporting the traditional style of pure mathematics research paper through the use of computation as part of the proof of theorems; in the construction of examples and counter examples; and in algorithms research. As they are drawn from published papers they reflect what is documented in archival publications - much is omitted. Notice first that it would be unusual to have a proof that consisted entirely of a GAP computation; such a proof would probably not be considered deep enough to warrant journal publication, unless it was given context as part of a larger body of work which was felt to be significant. Thus in our sample one paper is devoted essentially entirely to computations of 9 pages of tables together with a narrative explanation of them: Harvey and Rayhaun build evidence for a connection between a particular modular form occurring in the work of Field’s medal winner Borchers, and the representation theory of the Thomson sporadic simple group. This is related to remarkable results linking sporadic simple groups, modular forms and conformal field theory, popularly labelled “Moonshine”, which came about when Mackay observed in 1979 a connection between the “Monster” simple group and a certain modular form, through observing the role of the number 196883 in both.

However, sampling published papers to spot usage of GAP is also misleading, as it does not reflect significant and deep use of systems such as GAP in the process of doing mathematics, of exploring patterns, ideas and concepts, playing with examples and formulating and testing conjectures. We have heard, anecdotally, of mathematicians spending several months on calculations where the use of GAP was not even mentioned in the final paper: in one case a lengthy calculation involving the number 7 in GAP then informed published hand calculations which mimicked the computer calculation but with the 7 replaced by the variable p throughout; in another a lengthy computer search found a counterexample to a conjecture that could be readily described and shown to have the required properties without mentioning the computer search; in a third lengthy calculations in GAP were carried out to build evidence for a series of conjectures before time and effort was invested in a hand proof. Evidence for the use of computer algebra (though not GAP) in developing a proof can be drawn from the *poly-math 8* project where several participants were using Maple, Mathematica and SAGE for experiment, calculation and search as a matter of course, alongside the main argument, from time to time reporting the evidence they had found, and comparing their results. The crowdsourcing approach clarified computational techniques and apparent variations in results, and provided added confidence in the approach.

The use of computer support in this way is in line with the manifesto of “Experimental Mathematics”, ably presented in a series of books and papers by Bailey, Borwein and others. They articulate the possible uses of symbolic and numeric computation as:

- (a) Gaining insight and intuition;
- (b) Visualizing math principles;
- (c) Discovering new relationships;
- (d) Testing and especially falsifying conjectures;
- (e) Exploring a possible result to see if it merits formal proof;
- (f) Suggesting approaches for formal proof;
- (g) Computing replacing lengthy hand derivations;
- (h) Confirming analytically derived results.

The analysis of GAP papers above is consistent with the remark that while (a), (d) and (g) might all appear in current published papers, the rest are more likely to happen in the more speculative stage of the development of a proof. Bailey and Borwein argue that we need to think carefully about whether to allow a computation to be considered directly as a proof, and how to establish new standards for it to take its place in the literature. They go beyond the use of computation in support of traditional proof methodologies to assert:

Robust, concrete and abstract, mathematical computation and inference on the scale now becoming possible should change the discourse about many matters mathematical. These include: what mathematics is, how we know something, how we persuade each other, what suffices as a proof, the infinite, mathematical discovery or invention, and other such issues.

While these are all worthy of debate, and indeed along with *polymath*, HOTT, and the work of Gonthier and Hales are stimulating increasing discussion in blogs and other scientific commentary, it is not clear that the practice of mathematics, as evidenced by mathematics publications, is yet changing. For example, a glance at the Volume 24, Issue 1, of the Journal of Experimental Mathematics, published in January 2015 finds 12 papers. All follow the standard mathematical style of presentation with theorems, proofs, examples and so on. Of these, 4 use computation as in (d) above - for testing and especially falsifying conjectures, by exhibiting a witness found through calculation or search. A further 5 are of type (a) and use computation to numerically evaluate or estimate a function, and hence conjecture an exact algebraic formula for it. Two start with computational experiments which stimulate a conjecture that all elements of a certain finite class of finite objects have a certain property, and then prove it by (computational) exhaustion: so these may be described as (a) followed by (d). One presents an algorithm plus a running example, so perhaps also (a). It would appear that all exhibit (g).

6 Learning from the Everyday

Gowers, Tao, Villani and Wiles are extraordinary mathematicians, which makes their reflections of the process of doing mathematics both fascinating and atypical. The case studies described above allow us to look at the everyday activities of more ordinary mathematicians. It also allows us to draw a number of conclusions related to the practice of mathematics, attitudes to innovations, and the possible deployment of computational logic systems.

Stumbling Around in the Dark. By looking at how mathematics is done - or as Tao puts it ‘how the sausage is made’ - we get a more detailed view of Wiles’s ‘stumbling around in the dark’. Our examples highlight the role of conjectures, concepts and examples in creating a proof. Interestingly, they provide an evidence base to challenge Lakatos’s account of the development of proofs. To simplify somewhat, Lakatos presents a view of mathematical practice in which conjectures are subject to challenge through exhibiting examples, leading to modification of the hypothesis, or the concepts underlying it, with all the while progress towards a proof (of something, if not the original hypothesis) being maintained. While we certainly see this process at work in *polymath*, *minipolymath* and *mathoverflow*, this description suggests an all too tidy a view of the world, as we also see lines of enquiry abandoned because people get stuck, or make mistakes, or spot what might be a fruitful approach but lack the immediate resources of time or talent to address it, or judge that other activities are more worthwhile. Villani’s playful account of the development of a proof is particularly insightful about this aspect of research, and we also observe numerous paths not taken or dismissed in the *polymath* and *minipolymath* problems. Likewise many of the *mathoverflow* questions demonstrate a general wish to understand a particular phenomenon or example, or find out what others know about it. rather than asking a precise question about its properties.

Examples and Computation. Our case studies exhibit a variety of ways in which examples are used: straightforwardly as part of an existence proof; in the Lakatosian sense of a way to test and moderate hypotheses; to explain or clarify; or as a way of exploring what might be true or provable. Examples play an interesting role in collaborative endeavours: since the same example or phenomenon may occur in different areas of mathematics with different descriptions, sharing or asking for an example in *mathoverflow* may open up connections, or shed new light on a problem, or allow rapid interaction through a new researcher trying to understand how an unfamiliar concept applies to their own favourite family of examples. Computational methods allow construction, exploration, and retrieval of a much greater range of examples, and such techniques appear to be absorbed into the standard literature without comment, despite the well known limitations and non-reproducibility of computer algebra calculations.

Crowdsourcing, Leadership and the Strategic View. After a few iterations the *polymath* projects evolved to having a leader (most have been led by Tao) who took responsibility for overall guidance of the approach, drawing together

the threads every so often to write a long blog post (with a view to it being part of the published paper), and setting the discussion on a new path. Perhaps most striking, and worthy of further study, is the strategic decisions that are made about which route to pursue in a complex landscape of possible proofs. While these are the result of the intuition and insight of extraordinary mathematicians, when the participants comment on this, we find judgements informed by what has worked in the past in similar situations, assessments of the relative difficulties or of the various approaches, or the likelihood that the approach will lead to something fruitful, even if it is not likely to solve the whole problem. Frenkel and Villani give similar insights in their books, with a frequent metaphor being that of searching for the proof as a journey, and the final proof as a road-map for the next explorer. Marcus du Sautoy, writes on proof as narrative: “The proof is the story of the trek and the map charting the coordinates of that journey”. To continue the metaphor, collaboration enables new ways of exploration, to draw on different skills, and, crucially to share risks in a way that can make participants more adventurous in what they try out. In a section of *polymath* 8a for example, several participants are experimenting in an adventurous way with different computer algebra systems, and are joined by an expert in Maple who is able to transform and integrate the informal ideas and make rapid progress.

Institutional Factors in Innovation. All of our case studies show, in different ways, innovations in the practice of mathematics, through the use of machines to support collaboration, knowledge sharing, or calculations in support of a proof. What is noteworthy is that, however innovative the process, the outcomes of the activity remain unchanged: traditional papers in a traditional format in traditional journals, albeit with some of the elements executed by a machine. The reasons for this appear not to be any innate superiority of the format, indeed plenty have argued perceptively and plausibly for change, but the external drivers on research mathematicians. Research mathematicians are almost exclusively employed in the university system, either in the developed world, or in organisations in the developing world who are adopting similar norms and mechanisms, and are driven by the need to gather traditional indicators of esteem and recognition. The leaders of the field, such as Gowers, Tao, and Wiles, are perhaps best placed to resist these drivers, but are likewise aware of the pressures on younger colleagues - as evidence the discussions about authorship in *polymath*, and the advice not to spend too much time on it before tenure. Such pressures are active in other ways - for example publishing a so-called ‘informalisation’ of a formal proof in Homotopy Type theory in the high profile LICS conference, or in shaping decisions about how to spend ones time, so that, for example, the tactical goal of getting a paper written over a summer before teaching starts in the fall trumps loftier concerns.

Finally, what does this tell us about computational logic? We have described the “stumbling around in the dark” that currently seems a inevitable part of developing a proof, using as evidence the traces left by participants in collaborative activities on the web, and users of computer group theory systems. We have stressed the importance of a strategic view of proof, and the diversity and

sharing of risk provided by collaboration. While we have not yet studied this in detail, we note no evidence that the same is not true of developing large machine proofs. Understanding the collaborative development of human proofs should help us understand the collaborative development of machine proofs as well, and the best way to combine the two: Obua and Fleuriot's ProofPeer is making a start. Vladimir Voevodsky argues that computer proof will lead to a flowering of collaboration, as it enables trust between participants, who can rely on the machine to check each others work, and hence enables participants to take more risks, leading to much greater impact for activities like *polymath*.

At the same time, we see increasing recognition of the power of machine proof for mathematics: the work of Gonthier, Hales and Voevodsky to the fore. A recent triumph for SAT solving in mathematics was the discovery by Konev and Lisitsa in 2014 of a sequence of length 1160 giving the best possible bound on a solution to the Erdos discrepancy problem, resolving a question that had been partially solved in an earlier *polymath* discussion, which found a bound of 1124, which Gowers and others believed was best possible. Konev and Lisitsa write "The negative witness, that is, the DRUP unsatisfiability certificate, is probably one of longest proofs of a non-trivial mathematical result ever produced. Its gigantic size is comparable, for example, with the size of the whole Wikipedia, so one may have doubts about to which degree this can be accepted as a proof of a mathematical statement." It is an indication of how attitudes to computer proof have evolved since the more negative comments and concerns reported by Mackenzie 20 years before, that Gowers responded on his blog that "I personally am relaxed about huge computer proofs like this. It is conceivable that the authors made a mistake somewhere, but that is true of conventional proofs as well."

Our review of papers in group theory showed that they often contain significant amounts of detailed symbolic hand calculation of the kind that it would be straightforward to carry out in a proof-assistant, even though this is not current practice. Likewise machine assistance would surely confer some advantages in organising proofs that rely on complicated "minimum counterexample" arguments, a common pattern when considering finite simple groups. Similarly machine "book-keeping" would help in handling elaborate case-splits, as often occur in proofs of results about groups of order p^r , for all, or for all sufficiently large, primes, where the behaviour of different residue classes of $r \bmod p$ need to be considered. As Vaughan-Lee writes of the work mentioned above "all the proofs are traditional "hand" proofs, albeit with machine assistance with linear algebra and with adding, multiplying, and factoring polynomials. However the proofs involve a case by case analysis of hundreds of different cases, and although most of the cases are straightforward enough it is virtually impossible to avoid the occasional slip or transcription error." Since use of GAP is now accepted and routine in such papers, it is hard to see why use of a proof assistant could not be also.

In a panel discussion at the 2014 ceremonies for the Breakthrough Prize, the winners Simon Donaldson, Maxim Kontsevich, Jacob Lurie, Terence Tao

and Richard Taylor addressed computer proof in various ways: asking for better search facilities (Tao), wondering if “Perhaps at some point we will write our papers not in LaTeX but instead directly in some formal mathematics system” (Tao), and remarking “I would like to see a computer proof verification system with an improved user interface, something that doesn’t require 100 times as much time as to write down the proof. Can we expect, say in 25 years, widespread adoption of computer verified proofs?” (Lurie). Several speakers pointed to the length of time it can take for humans to be certain that a complex proof is true, and Kontsevich pointed out that “The refinement and cleaning up of earlier, more complicated stories is an important and undervalued contribution in mathematics.”. This last point chimes with an observation made by Steingart on the classification of finite simple groups: the concern that the protagonists had that, with the leading figures in the field growing older, and few new recruits as other areas now seemed more exciting, the skills needed to understand these complex proofs and fix, if necessary, any local errors were being lost, and the proof risked being ‘uninvented’. Perhaps ensuring that mathematics, once invented with such difficulty, does not become uninvented again, and that we don’t forget how to read the map. is the greatest contribution computational logic can make to the field.

Acknowledgements. Ursula Martin acknowledges EPSRC support from EP/K040251. This essay acknowledges with thanks a continuing collaboration with Alison Pease, and incorporates material from two workshop papers which we wrote together.

Further Reading

Barany, M., Mackenzie, D.: Chalk: Materials and Concepts in Mathematics Research. Representation in Scientific Practice Revisited. MIT Press (2014)

Frenkel, E.: Love and Math: The Heart of Hidden Reality. Basic Books (2014)

Gorenstein, D.: Finite Simple Groups: An Introduction to their Classification. Plenum Press, New York (1982)

Hadamard, J.: The Psychology of Invention in the Mathematical Field. Princeton (1954)

Mackenzie, D.: Mechanizing Proof: Computing, Risk, and Trust. MIT Press (2001)

Steingart, A.: A group theory of group theory: collaborative mathematics and the ‘uninvention’ of a 1000-page proof. Soc. Stud. Sci. **42**, 185–213 (2014)

Villani, C.: Birth of a Theorem. Random House (2015)

Wiles, A.: Transcript of interview on PBS. www.pbs.org/wgbh/nova/physics/andrew-wiles-fermat.html

The Power of Collaboration: *polymath*

Gowers, T., Nielsen, M.: Massively collaborative mathematics. *Nature* **461**, 879–881 (2009)

The *polymath* Blog. polymathprojects.org

The *polymath* wiki. michaelnielsen.org/polymath1

“Is massively collaborative mathematics possible?”, Gowers’s Weblog, gowers.wordpress.com/2009/01/27/is-massively-collaborative-mathematics-possible

Nielsen, M.: Reinventing Discovery. Princeton (2012)

D H J Polymath: The ‘bounded gaps between primes’ polymath project: a retrospective analysis’. *Newslett. Eur. Math. Soc.* **94**, 13–23

Examples, Conjectures, Concepts and Proofs: *minipolymath*

Mini-polymath3 discussion. terrytao.wordpress.com/2011/07/19/mini-polymath3-discussion-thread/

Minipolymath3 project. polymathprojects.org/2011/07/19/minipolymath3-project-2011-imo

Martin, U., Pease, A.: Seventy four minutes of mathematics: an analysis of the third Mini-Polymath project. In: *Proceedings of AISB/IACAP 2012, Symposium on Mathematical Practice and Cognition II* (2012)

Questions and Answers: *mathoverflow*

Martin, U., Pease, A.: What does mathoverflow tell us about the production of mathematics? arxiv.org/abs/1305.0904

Tausczik, Y.R., Pennebaker, J.W.: Participation in an online mathematics community: differentiating motivations to add. In: *Proceedings CSCW 2012*, pp. 207–216. ACM (2012)

Mendes Rodrigues, E., Milic-Frayling, N.: Socializing or knowledge sharing?: Characterizing social intent in community question answering. In: *Proceedings CIKM 2009*, pp. 1127–1136. ACM (2009)

Everyday Calculation: GAP

Vaughan-Lee, M.: Groups of order p^8 and exponent p . *Int. J. Group Theor.* Available Online from 28 June 2014

Bailey, D.H., Borwein, J.M.: Exploratory experimentation and computation. *Not. Am. Math. Soc.* **58**, 1410–1419 (2011)

GAP – Groups, Algorithms, and Programming, Version 4.7.7. The GAP Group (2015). www.gap-system.org

Harvey, J.A., Rayhaun, B.C.: Traces of Singular Moduli and Moonshine for the Thompson Group. arxiv.org/abs/1504.08179

Conclusions

Donaldson, S., Kontsevich, M., Lurie, J., Tao, T., Taylor, R.: Panel discussion at the 2014 award of the Breakthrough Prize. experimentalmath.info/blog/2014/11/breakthrough-prize-recipients-give-math-seminar-talks/

du Sautoy, M.: How mathematicians are storytellers and numbers are the characters. www.theguardian.com/books/2015/jan/23/mathematicians-storytellers-numbers-characters-marcus-du-sautoy

Lakatos, I.: Proofs and Refutations. CUP, Cambridge (1976)

Obua, S., Fleuriot, J., Scott, P., Aspinall, D.: ProofPeer: Collaborative Theorem Proving. arxiv.org/abs/1404.6186

Automated Deduction - CADE-25
25th International Conference on Automated
Deduction, Berlin, Germany, August 1-7, 2015,
Proceedings
Feltz, A.P.; Middeldorp, A. (Eds.)
2015, XXVIII, 640 p. 93 illus., Softcover
ISBN: 978-3-319-21400-9