The social machine of mathematics

Ursula Martin, Queen Mary, University of London
Ursula.Martin@eecs.qmul.ac.uk

"There is no ... mathematician so expert ... as to place entire confidence in his proof immediately on his discovery of it... Every time he runs over his proofs his confidence encreases; but still more by the approbation of his friends." *David Hume, 1739*

"Computers can ... create abstract social machines on the Web: processes in which the people do the creative work and the machine does the administration...The stage is set for an evolutionary growth of new social engines." *Tim Berners-Lee, 1999*

"... 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." *Andrew Wiles, 2000, on proving Fermat's theorem*

"Who would have guessed that the working record of a mathematical project would read like a thriller?" *Tim Gowers/Michael Nielson, on collaborative online mathematics, Nature, 2009*

"This is really the End." *Georges Gonthier, 2012 completes his 6 year formal verification of a major 255 page result in algebra, the odd-order-theorem*

For centuries, the highest level of mathematics has been seen as an isolated creative activity, to produce a proof for review and acceptance by research peers. Mathematics is now at a remarkable inflexion point, with new technology radically extending the power and limits of individuals. "Crowdsourcing" pulls together diverse experts to solve problems; symbolic computation tackles huge routine calculations; and computers, using programs designed to verify hardware, check proofs that are just too long and complicated for any human to comprehend.

Yet these techniques are currently used in stand-alone fashion, lacking integration with each other or with human creativity or fallibility. Social machines are new paradigm, identified by Berners-Lee [21], for viewing a combination of people and computers as a single problem-solving entity.

What if we developed a new vision, changing the way people do mathematics, and transforming the reach, pace, and impact of mathematics research, through creating a mathematics social machine — a combination of people, computers, and archives to create and apply mathematics?

Thus, for example, an industry researcher wanting to design a network with specific properties could quickly access diverse research skills and research; explore hypotheses; discuss possible solutions; obtain surety of correctness to a desired level; and create new mathematics that individual effort might never imagine or verify. Seamlessly integrated "under the hood" might be a mixture of diverse people and machines, formal and informal approaches, old and new mathematics, experiment and proof.

Much is known about the relevant ICT technologies:

- Collaborating: crowdsourced and open innovation
- Creating: AI for creativity, analogy and discovery
- Calculating: numeric and symbolic computation
- Verifying: formalization, reasoning and proof
- Sharing: knowledge management and interfaces

The obstacles to realising the vision seem to be not advances in any one of these domains, but rather:

- We do not have a high level understanding of the production of mathematics by people and machines, integrating the current diverse research approaches
- There is no shared view among the diverse research and user communities of what is and might be possible or desirable

In this note we sketch what we might do to address these challenges.

Background

Mathematics reach, pace and impact Mathematics and theoretical computer science, research underpins modern programming languages, secure systems, and the Web. Advances depend on hard foundational mathematics, and draw on newer areas such as statistics and dynamical systems, alongside traditional combinatorics and logic, supplemented by simulation and experiment. This potential *reach* of mathematics is increasing, thus increasing the challenge to researchers of deploying the right combinations of techniques within and beyond their own specialism to solve increasingly hard and broad problems, which are not stated in isolation but require an integrated approach. The rapid *pace* of technology creates key opportunities, if mathematicians are able to collaborate with each other, and with other computing disciplines, both academic and industry, to produce timely results. Increasing pace and reach has the potential to increase *impact*, if potential users of research can find the researchers and the research they need. Sometimes it can be easier to write a new paper than to find old results: the past 10 years saw nearly 200K papers of relevance to theoretical computer science [2].

Social machines Social machines combine people and computers for emergent and collective problem-solving. Current examples include Google, Wikipedia and Galaxy Zoo, providing platforms for innovation, discovery, and commercial opportunity [34]. Future more ambitious social machines will combine deep social involvement and sophisticated automation [21], and are now the subject of major research. This approach builds on e-Science work such as Goble's myExperiment [15], a collaborative research space for scientific workflow management and experiment: however such systems do not address mathematics.

Mathematics and social machines The production of mathematics provides an important, timely and exciting challenge for social machines research — with a variety of approaches to combining people and machines. In the past few years, systems for unstructured collaboration developed by researchers themselves have had a powerful impact: we call such systems *social mathematics*.

- in the summer of 2010 a paper was released plausibly claiming to prove one of the major challenges of theoretical computer science, that $P \neq NP$. It was withdrawn after rapid analysis coordinated by senior scientist-bloggers
- *polymath* collaborative proofs, a new idea led by Gowers, use a wiki for collaboration among scientists from different backgrounds and have led to major advances [18]
- discussion fora, including new ideas such as user ratings for finding the right expert, allow rapid informal interaction and problem solving; in three years *mathoverflow.net* has hosted 27,000 conversations
- the widely used "Online Encyclopaedia of Integer Sequences" (OEIS) invokes subtle pattern matching against over 200K user-provided sequences on a few digits of input: so for example (3 1 4 1) returns π [7]
- the *arXiv* holds around 750K preprints in computer science, mathematics etc.. By providing open access ahead of journal submission, it has markedly increased the speed of refereeing, widely identified as a bottleneck to the pace of research [31]
- Innocentive [22], a site hosting open innovation and crowdsourcing challenges, has hosted around 1,500 challenges with a 57% success rate, of which around 10% were tagged as mathematics or ICT.

All can be viewed as social machines — for example OEIS involves users with queries or proposed new entries; volunteers curating the system; governance and funding mechanisms; as well as a database, matching engine and web interface.

The social element The social is crucial in the production of mathematics. Mackenzie's sociological study of proof [25] confirmed Lakatos's analysis of the role of error [23], and Hume's assessment nearly 300 years earlier of the social nature of proof [10]. Williams's notion of technological "artefacts" matches the way in which mathematical objects mutate as ideas are developed [41].

The work of cognitive scientists, sociologists, philosophers and the narrative accounts of mathematicians themselves, highlight the paradoxical nature of mathematical practice — while the goal of mathematics is to discover mathematical truths justified by rigorous argument, mathematical discovery involves "soft" aspects such as creativity, informal argument, error and analogy.

Collaborative systems such as *polymath* contribute to mathematics research, and also provide a rich evidence base for further understanding of mathematical practice. Our analysis of a *polymath* proof [35] found only 47% of the conversational "turns" were proof steps, with the rest being made up of conjectures, concept formation and the like.

At a recent learned society event organised by the proposer [1], leading mathematicians flagged the importance of collaborative systems that "think like a mathematician", handle unstructured approaches such as the use of "sloppy" natural language, and the exchange of informal knowledge and intuition not recorded in papers, and engage diverse researchers in creative problem solving.

Yet if the mathematics social machine is to realise its potential, and scale up to large proofs, it will also need formal approaches. Verification through formalisation and proof, supported by decades of academic and industry research into theorem provers, is achieving remarkable breakthroughs, and providing rich archives of material for possible re-use:

- on 20th September 2012 Georges Gonthier announced that after six years effort he had completed a formalisation in the Coq theorem prover of one of the most important and longest proofs of 20th century algebra, the 255 page odd-order theorem [3]
- mathematician Tom Hales has almost completed a ten-year formalisation of his proof of the Kepler conjecture, using several theorem provers to confirm his major 1998 paper [19]
- hardware and software verification to ensure error-free systems is a major endeavour in companies like Intel and Microsoft [20], as well as supporting a number of specialist SMEs.

Other likely elements of a mathematics social machine would include the following, all currently major research activities in their own right:

- a variety of AI and cognitive science inspired approaches to "soft" aspects such as creativity, analogy and concept formation [13]. For example, mathematicians often mention the importance of error for creativity [1]: this has also stimulated Bundy's recent AI work on ontology repair [26]
- symbolic and numeric computation, and associated data, provided by commercial systems such as Matlab and Maple, or research packages such as GAP: all already engage strongly with e-science
- digitised mathematical archives, using MKM, for example to support search, re-use and executable papers [24]. The National Academy of Sciences have just announced a major initiative [33].
- interfaces: people to machine, natural language to mathematics, and software to software

Capitalising on the substantial research underlying these achievements will inform thinking about the design space for mathematics social machines, for example:

- precise versus loose queries and knowledge
- human versus machine creativity
- specialist/niche versus general users
- logical validity versus cognitive appeal for output
- formal versus natural language for interaction
- generating versus checking conjectures or proofs
- formal versus informal proof
- "evolution" versus "revolution" for new systems
- governance, funding and longevity

Exploring these in the framework of social machines will include matters such as:

- *Designing social computations* Social machine models [21] view users as "entities" (cf agents or peers) and allow consideration of social interaction, enactment across the network, engagement and incentivisation, and methods of software composition that take into account evolving social aggregation. For mathematics this has far reaching implications — handling known patterns of practice, and enabling others as yet unimagined, as well as handling issues such as error and uncertainty, and variations in user beliefs.
- *Accessing data and information* Semantic web technology enables databases supporting provenance, annotation, citation and sophisticated search. Mathematics data includes papers, records of mathematical objects from systems such as Maple, and scripts from theorem provers. There has been considerable research in MKM [24], but current social mathematics systems have little such support. Yet effective search, mining and data re-use would transform both theoretical computer science research and commercial verification. Research questions are both technical, for example how to ensure annotation remains timely and correct, and social, for example many *mathoverflow* responses cite published work, raising the issue of why users prefer asking to searching.
- *Accountability, provenance and trust* Participants in social machines need to be able to trust the processes and data they engage with and share. Key concepts are *provenance*, knowing how data and results have been obtained, which contributes to *accountability*, ensuring that the source of any breakdown in trust can be identified and mitigated [40]. Current social mathematics systems are remarkably open — for example posting drafts on the *arXiv* ahead of journal submission is reported as speeding up refereeing and reducing priority disputes [1]. Trusting mathematical results requires considering provenance of the proof, a major issue in assessing the balance between formal and informal proofs, and the basis for research into proof certificates [30]. Privacy and trust are significant for commercial or government work, where revealing even broad interests may already be a security concern.
- *Interactions among people, machines and data* Interactions among people, machines and data are core to social machines, which have the potential to support novel forms of interaction and workflow which go beyond current practice, a focus of current social machine research [21]. Social mathematics shows a variety of communities, interactions and purposes — looking for information, solving problems, clarifying information and so on [35] — displaying much broader interactions than those supported by typical mathematical software. Lakatos identifies mathematical "moves" (eg responses to counterexamples) that are examples of mathematical workflow, and examining both *polymath* and the production of formal proofs has potential to reveal more [35]. In particular such workflows need to take account of informality and mistakes [12].

References

[1] http://events.inf.ed.ac.uk/sicsa-mcp/.
[2] Microsoft academic search, 27 June 2012.
[3] http://tinyurl.com/8hyegj2
[4] G Helleloid, U Martin The automorphism group of a finite p-group is almost always a p-group. *J Algebra* 312 (2007).
[5] U Martin et al. Sound and complete Floyd-Hoare logics. *ACM Trans Comput Log*, 11, 2009.
[6] The economic benefits of UK Mathematical Science Research, Deloitte/RCUK, to appear 2012

[7] OEIS Foundation Inc. (2011), The On-Line Encyclopedia of Integer Sequences, http://oeis.org.

[8] T Alpcan and T Basar. *Network Security: A Decision and Game Theoretic Approach* CUP, 2011

[9] D K Arrowsmith Fair sharing of resources in a supply network *Phys. Rev. E* 85 (2012) 46101–14

[10] D Hume *Treatise on human nature*

[11] U Martin et al Design verification for control engineering. In *IFM*, LNCS 2999 (2004) 21–35

[12] A Bundy et al Reasoning with Context in the Semantic Web, *J Web Semantics* 12 (2012)

[13] A Bundy, Automated theorem provers: a practical tool?, *Ann Math Artif Intell* 61 (2011) 3–14

[14] J Cheney et al Provenance as dependency analysis Math Struct Comp Sci 21(2011) 1301—37

[15] D De Roure, C Goble. Software design empowering scientists. *IEEE Software*, 26 (2009) 88–95

[16] U Martin et al Real number theorem proving in PVS, *Formal Aspects of Comp*, Online First 2012

[17] H Gottliebsen, T Kelsey and U Martin Hidden verification for computational mathematics *Journal of Symbolic Computation* 39 (2005) 539–567

[18] T Gowers and M Nielsen. Massively collaborative mathematics *Nature*, 461 (2009) 879–881

[19] T Hales et al. A revision of the proof of the Kepler Conjecture. *Discrete & Comp Geom* 44 (2010)

[20] J Harrison *Handbook of Practical Logic and Automated Reasoning* CUP, Cambridge, 2009

[21] J Hendler and T Berners-Lee From the Semantic Web to social machines: A research challenge for AI *Artificial Intelligence* 174 (2010) 156 – 161

[22] www.innocentive.com retrieved Oct 2012

[23] I Lakatos *Proofs and Refutations*. CUP 1976.

[24] M Kohlhase et al Semantics of OpenMath and MathML3 *Math in Comp Sci* to appear 2012

[25] D MacKenzie *Mechanizing proof: computing, risk, and trust* MIT Press, Cambridge, MA 2001.

[26] A Bundy et al, Dynamic first-order ontology repair, *J Semantic Web and Inf Systems* 3 (2007) 1–35

[27] P Malacaria, Security threats for looping constructs, *J Comp Security* 18 (2010) 191-228

[28] L Meagher et al, Flows of knowledge and influence *Research Evaluation* 17 (2008) 163-173

[29] J Mesirov. Accessible reproducible research. *Science*, 327 (2010) 415 – 416

[30] D Miller, Broad Spectrum Proof Certificates, *Cert Progs and Proofs*, LNCS 7086 (2011) 54–69

[31] J Crowley et al, Mathematics Journals: What May Change *Not AMS* 58 (2011) 1127–1130

[32] C. Neylon. Re-use as impact. In *altmetrics11, workshop at ACM Web Science Conference*, 2011.

[33] http://tinyurl.com/cwmtkxq

[34] M Nielsen. *Reinventing Discovery: The New Era of Networked Science*. Princeton, 2011.

[35] A Pease and U Martin. Seventy four minutes of math *Proc AISB Symp Math Practice* (2012) 19–29

[36] http://tinyurl.com/9napbjc

[37] S Riis Information flows, graphs and their guessing numbers. *Electr J Comb*, 14 (2007)

[38] SOCIAM: the theory and practice of social machines http://tinyurl.com/cmmwbcf

[39] V Stodden. Reproducibility in the Computational Sciences. MIT Paper 4773–10, 2010.

[40] D Weitzner et al, Information accountability, Commun. ACM, 51(2008) 82–87

[41] R Williams et al, From Artefacts to Infrastructures *Comp Supported Cooperative Work*, 21, 2012