A Little Rebellion Now and Then Is a Good Thing

Views on the Requirements Engineering Conference

Tony Gorschek

School of Computing
Blekinge Institute of Technology (BTH)
Karlskrona, Sweden
tony.gorschek@bth.se

Abstract— **A little rebellion now and then is a good thing.**

This short position statement describes my views on some of the challenges associated with many conferences, the Requirements Engineering Conference being among them. The main concepts are; the goals, as well as criteria for paper selection for the conference should be defined explicitly, and shared with the community. Industry involvement in the conference should be increased, but the focus of all tracks should be quality – what constitutes quality however needs to be defined and agreed on. Industrial validation of research results have to be more than an intention. Last but not least, how papers are presented and discussed needs to change, focusing on quality over quantity.

I. INTRODUCTION

Recently I had the pleasure of reading two position papers that advocated change. I am a big believer in renewal, and a little revolution (rebellion[1]) now and then is definitely a good thing. I found myself agreeing with both papers; the funny thing was that they were disparate, expressing (at least on the surface) two opposite directions and views. The paper by Wieringa et al. [1] focused on returning the RE conference to researchers, and removing the industry paper track completely. The two main points were that industry tracks have not been successful in attracting real (or good) industry contributions, or substantial industry presence. Their suggestions were that the RE conference could have a case study (empirical) track instead of an industry track, and to keep in touch with industry researchers should visit industry conferences.

The second paper by Mahaux and Mavin [2] focused on how to improve the (industrial) relevance of the RE conference, proposing a manifesto with significant implications on an outdated traditional conference model (and peer-review system). Their suggestions focused on moving away from the numbers game (H-index, acceptance rate), and towards a more collaborative (wiki-style) peer-peer review model that emphasized collaboration amongst contributors.

My summary does not do their positions justice, thus I highly recommend reading both papers and form your own opinion, but more importantly, propose ideas for solutions to the challenges we most certainly have. Weather you subscribe to one view, or think both are wrong, I think most can rally around the fact that we have challenges.

This short position paper expresses my own views on the central challenges facing requirements engineering (RE) research in the present (and past) decade as pertaining to its premiere forum, the Requirements Engineering Conference.

II. WHO AM I AND WHY SHOULD YOU LISTEN?

I did an engineering degree, then a BSc. in software engineering, quit school and started a company. I had to sell all my belongings to afford equipment and offices, and a couple of equally motivated friends/hackers joined me in starting what was to be (in our minds) the next Apple/Microsoft/Google – only better, bigger and cooler. We sat and coded, 80 hours a week, every week, for six months, sleeping in the office (no money for rent), before we dared show our idea to anyone. Showcasing our (MV)Product to capitalists and investors resulted in a lukewarm reception, to our great surprise. We shook it off and summized: 'people are idiots, don't they see what we have here?' Eventually we got some capital, and grew, continuing our frantic journey for a year.

We didn't do any requirements engineering, no design, no project management, no… anything really. We coded, prototyped (some), used a whiteboard to hash out problems (stealing pens from the university to save money), and released our product – and, oh yes – failed spectacularly.

Since then I have founded another five startups (three very successful ones), made money, worked as a senior executive consultant, done a MSc., done a PhD, become a professor, and pursued interests in both academia and industry. Today I generally sit on the other side, as an investor, and I have left the coding business to my betters, and now I focus on engineering processes and practices mostly.

Despite having experienced a lot, both in academia and industry, my most profound revelation is from my first (spectacular) failure. Knowing what your customers want and need will make or break you. To me, this is the core of requirements engineering (and product management for that matter). Somewhat simplified, the rest is about how to "best" specify, analyze, select, convey, validate, verify…[more]… the needs and wants.

---

[1] From "A Letter From Thomas Jefferson To James Madison"

This is how I got into the research game – I wanted to learn from others (and I should have read more in the early days before doing) - but also solve some problems I had seen for myself. (Weather I have is not for me to judge).

III. So, What Are the Challenges?

To be honest, there are many challenges; too many to do justice in a short position paper. But I will highlight items that I consider important, which are partly linked to the two (much better) position papers contributed by my colleagues.

A. Peer-Review as We Know It

I think the peer-review system, as it is intended to work, is a great idea. However, many times the realization of the idea is lacking. We have senior experts (read: researchers) that have arrived at their status through many ears of good work | networking | paying their dues | contribution | other. They generally constitute the "cream of the crop" in deciding what gets published and what does not. In conferences you see them as the program board (PB) and in journals as editors. Optimally they use reviews as input, and do a fair and impartial assessment, prioritize among contributions, and pick the best for publication. In the worst-case however personal opinions, preferences, and politics influence decisions.

Now, I have not been at many PB meetings, but the ones I have attended (e.g. RE and ESEC-FSE) have been professional and with a genuine intention to accept quality publications. That said, of course personal views and subjectivity exist among PB members, we are all humans, and these influence decisions. However, I think the problem goes deeper than that. We do not even agree on the basics. What is quality? What is relevance (and for whom)? What about rigor? Usefulness and Usability? Ask ten researchers to qualify these concepts and you will get twenty (different) answers, maybe more.

We did a review of publications in the Requirements Engineering journal and in the Empirical Software Engineering journal [3, 4] in 2009/10. The reading was rather grim. Most papers that published actual RE "solutions" were not rigorously written - to say the least (thus a reader could not judge what had been done, or exactly how), and the use of toy examples dominated what little validation there was. Right or wrong, this review was done using our definitions of the concepts. But at least we tried to define, and explain our use of, rigor and relevance.

I think the fundamental problem is in that we do not have one view (or even one combined complex view) of what we consider good, relevant, rigorous, useful or usable. The overwhelming opinion is that "we know it when we see it". This is not good enough.

IDEA 1: The steering group (SC) and the program board (PB) should sit down and discuss their views on the subject, formulate a detailed account of views, and make this public and available to the community. In short, define all relevant concepts and exemplify good and bad examples. I can volunteer some bad ones from my own papers if it helps.

Also, ask for input from the community – a sort of peer-review of the basis for the RE conference peer-reviews. If the views of the people "in charge" differ from the community at large this is opportunity for input, discussion and renewal.

Democracy is the worst form of government, except for everything else that has been tried[2]. I think the same can be said for the peer-review system. That does not mean that we should not aspire to perform peer-review to the best of our ability – and this requires openness and transparency in terms of what our judgment criteria are, and how we use them.

B. Industry Relevance, Relevance for Academics – or What about Relevance Pure and Simple

I don't know how many times I have had the "relevant for practice" discussion with fellow researchers and practitioners, and most times I get frustrated. The simple truth of the matter is that I get the feeling that there are two camps, the "real" research camp, where you do "advanced stuff", and the "other" camp where "very-close-to-industry" work is done. Some researchers do both, but honestly, most don't. How many papers have you read where a model/solution/other has been presented, exemplified on a toy example, and in the *future work section* validation in industry is promised as a next step. Then a year later additional papers show up that build on previous work, introducing more advanced concepts, building on the first paper, and yet again the future work section promises validation. And around we go.

I am not saying that research has to be validated or empirically grounded in all cases. What I am saying is that requirements engineering research is an engineering science[3] – and per definition concepts like scalability, usability and usefulness have to be taken seriously. If we as engineering scientists do not, who will? Consultants? I have a worked as a senior executive consultant for several years, and let me tell you, the name of the game is billable hours, not problem solving or transferring research to practice.

IDEA 2: What do you say about a 50-50 split? Half of the papers at an engineering conference (like RE) could be focused on empirical/validation studies (getting/using data from industry, transferring/testing out new concepts in industry, collecting challenged from industry and so on), and half could be about trying to solve problems/new solutions/other (and in this case validation does not have to be prioritized at all). Also, the validation needs to be properly designed, executed and reported - here rigor is central. However, it is very challenging to working with and in industry, thus a realistic view of what is good-enough has to be applied [6, 7].

Alternatively, as defined in IDEA 1, we transparently show that application of results in industry are not the focus of the conference through the criteria we apply for paper selection, giving other conferences the chance to fill the gap.

IDEA 2.5: Industry tracks (and I have worked with a couple) are a great idea. However, their practical realization has been lacking. Often the industry track ends up with lower

---

[2] Winston Churchill

[3] Definition: The discipline dealing with the art or science of applying scientific knowledge to practical problems

quality papers that are not nessecaraly empirical at all, and more than not written by researchers. Sometimes (and trust me, the PC chairs of the industry tracks work hard to get and accept quality contributions) it can even become the "didn't-really-make-the-research-track" track. So, what is the problem? Well, there are many, but two stand out. First, the industry track cannot be seen as the "poor cousin", either we are serious about the quality, but apply different criteria on what should be accepted (see IDEA 1), or we simply have the courage to cancel the track given too few quality contributions. This also means that if we subscribe to the 50-50 idea (see IDEA 2), then half of the resources and efforts of the conference should go into promoting, eliciting, and hunting for, industry track contributions.

In addition, very few practitioners write papers, so it is completely acceptable that researchers do, but they should be part of the 50%-empirical/validation. In practice it will be an empirical track (but we should call it Industry Track as it conveys what we need).

IDEA 3: Also, to be sure that the industry views are represented in decisions, why not invite one industry reviewer per paper (for both tracks) and have his/her opinion count as much as the two academics reviewing the paper?

Honestly, I think the industry tracks are improving, so are the contributions (and I don't say that just because I was a Industry PC co-chair in 2012). However, we cannot use the same base criteria to judge and select papers for the industry track as we do for the research track… and we must never, under any circumstances, see the industry track as "less". There have been great papers and good research presented at industry tracks, we just have to make sure we get more.

IDEA 3.5: How many years of combined research experience is represented in the SC, PB, and PC? Same question, but exchange "research experience" for industry experience. The RE conference is a research conference, but maybe, just maybe, some industry representatives should be present in the steering of the conference? I know I am suggesting something that will be very hard to accept, my experience is simply this; a heterogeneous group will often produce more overhead, but also better ideas.

C. The Numbers Game

When reading the two other position papers it became obvious that motivation (read: The reward systems) is at the core of many of the challenges. Sadly, the numbers game [5, 1] largely dictates the solution space of improving a conference, or a field of research for that matter. Citations, number of publications, and publication venue are important to most researchers. Amount of submissions, number of attendees, and acceptance rate are important to most conferences. Thus many of the noble concepts and good ideas that we aspire to can be in conflict with the reward systems we rely on (are forced to optimize against).

So, do we rebel and say 'to hell with it'? Well, to some extent maybe we should. Some have; medicine, physics and biology to mention a few, have moved to open access venues, often as a rebellion, and it has contributed to their renewal.

Other examples are break-out splinter conferences created out of the conviction of a few going against the current.

However, as long as we are punished (absence of reward) for doing what we consider the right thing, change will be slow. The trick is to find room to maneuver within the reward systems – which is a sort of rebellion in its own right.

IDEA 4: Accept fewer papers, but spend more time on each (I can imagine some reactions to this statement, but read on).

When at a conference I find very little real value in attending the paper tracks. Presentations are short and discussions are symbolic – I can confirm what Mahaux and Mavin [2] write. To be honest, sitting in on paper tracks I mostly read emails or work, listening with half an ear. Don't get me wrong, I want to meet and discuss research with peers, but you do that in-between/after the tracks, or during social events. As for getting an idea of what the papers are about, browsing the proceedings will get you just as far as listening to a presentation (in most cases).

The best model I have experienced was at REFSQ (before it became a "working conference"), where each paper was presented for ten-fifteen minutes, and then a discussion was held for thirty-forty minutes! Further, the discussions were led and moderated by two authors (of other papers) that were prepared and had read the paper in advance, preparing questions! The discussions were great, and the collaborative atmosphere even better. I hear the critique loud and clear: '*This format is for workshops and will not work in a conference.*' Well, I beg to differ. First, accept as many papers as you normally do (or more), for the proceedings, but select a smaller subset for actual presentation and discussion. If a PC/PB can choose what papers to accept, why can they not choose what papers would be most interesting to present and discuss? Or, why not have a vote, on-line, a month before the conference? Why not both?

This model would make the paper tracks interesting to attend, and at the same time not really impact researcher's ability to go to the conference (as they are published in the proceedings even if they do not present). Also, browsing the proceedings I can still find fellow researchers published in the proceedings and talk to them about their research – not presenting does not factor in.

IV. The End

First, I think that we need new and good ideas, renewal, and some measure of courage to succeed in making the RE conference the best and most relevant forum for research in the field. Today it is not. Being open, transparent, explicit, and listening to the community are pre-requisites. The courage part is realized in the willingness to change (and to some extent take risks). Why not try some of the new ideas presented in the position papers ([1, 2]). Even better, why not put all "ideas" up for a vote? Let the SC/PB vote count for 5 each - but the one winner each year is realized in the next year's conference.

I have the great honor of being the GC for RE'14 in beautiful Karlskrona, Sweden, next year (see www.RE14.org).

I would be happy to get the ball rolling and we can have the first live "vote for change".

As long as you play it safe you minimize the risk of sudden trauma - but that is a boring way to live your life. The RE conference is turning twenty-two, I think it can handle some risk in the pursuit of innovation and renewal, don't you?

BIO

Dr. Tony Gorschek is a Professor of Software Engineering at Blekinge Institute of Technology. His research commitments are both national and international ranging from conference organization (GC RE'14), PC work in premiere conferences (ICSE, ESEC/FSE), to serving on the editorial board for the premiere Springer journal Empirical Software Engineering.

Dr. Gorschek has over ten years industrial experience as a CTO, senior executive consultant and engineer, but also as chief architect and product manager. In addition he has built up six startups in fields ranging from logistics to internet based services and algorithmic stock trading.

Currently he manages his own consultancy company, works as a CTO, and serves on several boards in companies developing cutting edge technology and products. At present he supervises four PhD students and leads several research projects. As example can be mentioned a project with industry partners Axis and Volvo in the area of Requirements and Verification & Validation co-optimization. IBM, Ericsson and Qtema are partners in a project dealing with Large-scale software based innovation where an innovation candidate selection engine is under development. Daimler AG and Volvo, both, a luxury car manufacturers, are also industrial partners where he working with creation of models for optimized communication and lean and agile environments in large scale development of software intensive systems.

For more information see www.gorschek.com

REFERENCES

[1] R. Wieringa, P. van Eck, J. Mylopoulos, "Requirements Engineering Conferences: Wither Industry Tracks?", in proceedings of the 21st IEEE Intl. Requriements Engineering Conference, Rio, 2013.

[2] M. Mahaux and A. Mavin, "Towards <something else>: A New Paradigm for Applied Requirements Engineering Research", in proceedings of the 21st IEEE Intl. Requriements Engineering Conference, Rio, 2013.

[3] Ivarsson M., Gorschek T., "A Method for Evaluating Rigor and Industrial Relevance of Technology Evaluations", Empirical Software Engineering, Volume 16, Number 3, pp. 365-395, 2011.

[4] Ivarsson M., and Gorschek T., "Technology Transfer Decision Support in Requirements Engineering Research: A Systematic Review of REj", Requirements Engineering journal, Vol. 14, Number 3, 2009, pages 155-175.

[5] D. L. Parnas, 'Stop the numbers game', Communications of the ACM, vol. 50, no. 11, pp. 19–21, 2007.

[6] Gorschek T, Garre P, Larsson L, Wohlin C, "A Model for Technology Transfer in Practice", IEEE Software 23(6), 2006, pp. 88-95.

[7] Wohlin C., Aurum A., Angelis L., Phillips L., Dittrich Y., Gorschek T., Grahn H., Henningsson K., Kågström S., Low G., Rovegård P., Tomaszewski P., van Toorn C., and Winter J., "The Success Factors Powering Industry-Academia Collaboration", IEEE Software, vol. 29 no. 2, 2012, pp. 67-73.