Requirements Engineering Conferences: Wither Industry Tracks?

Roel Wieringa, Pascal van Eck
University of Twente
The Netherlands
(r.j.wieringa | p.vaneck)@utwente.nl

John Mylopoulos
University of Trento
Italy
jm@cs.toronto.edu

Abstract—This position paper argues that industry tracks have no place in any research conference. Instead, a research conference should always have room for industrial case studies, evaluated according to criteria for empirical research. Such case studies would not be acceptable at a practitioners' industrial conference, just as papers presented at such conferences would not be acceptable at research conferences. It follows as corollary that if researchers want to become familiar with problems and solutions of RE practice, they should visit industrial conferences.

Index Terms—Industry track, case study track, research conference, practitioner conference

I. Introduction

Conferences in software and requirements engineering organize industry tracks intended to bring together researchers and practitioners so that they can exchange ideas about problems and solutions with relevance to practice. This short discussion paper argues that these tracks should be removed from these conferences. We start with our two recommendations:

First, rather than attempting to attract practitioners to research conferences, it seems more fruitful for researchers to visit real industry conferences to learn about the concerns of practitioners and about solutions in the real world. Researchers should also become members of professional networks where practitioners exchange experiences and discuss possible, practical solutions. Relevance of research is founded on selection of relevant problems, and the best way to learn about relevant problems is for researchers to meet practitioners on their own ground.

Second, research conferences should stop organizing industry tracks, and start organizing scientific case study tracks. Industry tracks fail to achieve the purpose of bringing together researchers and practitioners. All too often, industry tracks merely consist of submissions forwarded by the program committee of the corresponding research track, because these submissions were deemed to not be "good enough" for a research track. Moreover, industry tracks generally do not attract real practitioners, but researchers interested in practice, such as researchers in industrial laboratories. Rather than organize industry tracks we should organize industrial case study tracks, where research papers are presented about methodologically sound industrial case studies. We have not yet realized the full potential of industrial case studies at the RE conference.

We will argue for these recommendations by describing the differences between research and industry conferences, illustrating this with an example of an industrial conference. We then describe in more detail what criteria a scientific case study should satisfy, and then argue in more detail that industry tracks should be replaced by case study tracks. We then turn to the history of the RE conference and trace the paper categories as they have been defined for the conference since the mid-1990s. This analysis shows that the industrial track of the conference has contained the industrial problems and experience paper category that we think should be removed, but also at times has contained the case study paper category, that we think should be moved to the research track.

II. The Difference Between Research and Industry Conferences

We should keep it simple: Industry papers are submitted and presented at industry conferences, research papers at research conferences. The gap in expectations, audience and presentation between these two kinds of conferences is so large that we should not attempt to mix them. Research conferences are events where researchers publish peer-reviewed research papers and meet to discuss ongoing work, explore opportunities for new project funding and cooperation, and where talented young researchers find new alliances with senior researchers. Industry conferences, by contrast, are events where enterprises and independent consultants meet to present their latest products and services, exchange their experiences and expectations, meet with other practitioners to explore opportunities for new business, and where practitioners find new employment opportunities. All of this contributes to the generation of new transactions in a commercial market, and also reduces the cost of these transactions.

The means for researchers and practitioners to achieve their goals are similar: publication in the form of a presentation and/or paper, and networking. However, the goals to be achieved by these means are totally different: the researcher's goals are to produce a peer reviewed publication; the practitioners goals are versus to generate transactions in a commercial market. The quality criteria for publications and presentations, relevant for these goals, are likewise disjoint.

A failed contribution for one of these conferences is *not* a possible contribution for the other. To check this, send a paper

rejected by a research conference to an industrial conference and see what happens.

Researchers who take the step to commercialize their research product appreciate the large difference between these two kinds of conferences. Once they are working with their own start-up company, the decision to attend an industry conference will be guided by the commercial opportunities present at such a conference, compared to alternative ways of spending their hours for the company. Research conferences will not appear on their radar as a means to achieve their commercial goals.

III. EXAMPLE

To illustrate these differences, we present as an example of a successful industry conference the two-day Dutch National Architecture Conference for the Digital World (with the Dutch acronym "LAC")[1] that has been held for almost 15 years. LAC attracts about 600 attendants annually, most of them senior enterprise architects, consultants, or IT managers, both from vendors of IT services as well as from user organizations such as bank/insurance companies, or government agencies. LAC is organized jointly by a commercial event organizer and an independent program committee (PC, of which the second author is a member) under auspices of the Dutch Architecture Forum for the Digital World, called NAF, a non-profit association of about 100 vendors and users of IT services and knowledge institutes such as universities, of which the first author has been a board member[2]. LAC is funded partly through the conference fee and partly through sponsorship, mostly by vendors of IT services. Sponsorship does not include a slot in the conference program.

More than half a year before the conference, the LAC PC distributes a call for contributions in which the Dutch architecture community is invited to submit proposals for tracks, individual presentations, tutorials or workshops. The call for contributions lists the theme of the year, as well as a list of suggested topics. Interested contributors submit a 1-2 page description of a proposed track (a set of five presentations around a common topic) or of a single presentation. There is no peer review process. Instead, the PC assesses these proposals on their perceived value of the proposed topic and content for the LAC audience. For one thing, this means no marketing talk: the unique selling point of a vendor and sponsorship-independent conference is that vendor pitches of a particular product or service are avoided.

Over the past few years, the PC accepted about half of all proposed contributions. Many of these contributions are presentations of what can be called "consolidated experience" of the presenter, codified in a set of lessons learnt or a conceptual framework. What differentiates these contributions from research papers is that they are personal: the contribution presents the personal opinion and perception of the presenter, not the result of a methodologically sound, peer-reviewed

empirical investigation. The value of these contributions for the audience of professionals is that it gives them new ideas for how to achieve practical goals.

If such a paper would be submitted to a research conference, it would be regarded as a presentation of subjective opinions biased to one company. For a potential customer, this would be useful information because it allows them to decide whether or not they can use this product or practice in their particular business context. But for a researcher, this is not interesting as there is no research question, no discussion of related work, and no justification of generalizability. So these papers would not be accepted at a research conference.

IV. INDUSTRIAL CASE STUDIES

In our field, case studies are by definition industry case studies: A scientific study of a particular case in practice. What should a case study look like? There are different guidelines on this, but a consensus seems to be emerging [1], [2], [3], [4]. First, all case studies should present and motivate their conceptual framework, identify the case to be studied, list their research questions, and summarize the current state of knowledge to which the case study aims to contribute.

Second, measurement instruments should be explained, such as questionnaires, interviews, focus groups, primary documents analyzed, etc. and the procedures to objectively interpret data needs to be explained.

Third, the availability of data for other researchers should be indicated (assuming that confidential information is anonymized or removed), observations should be summarized in numerical and/or graphical form, and explored on their consequences, and one or more possible explanations in terms of relevant theories need to be provided. And as any good research study, an industry case study should end with new research questions to be studied in later research.

Fourth, a case study should make clear whether it is an *observational* case study, in which the researcher tried to minimize their influence on the case, or an *action* case study, where the researcher intervenes in the case and studies the result. The list of items just given applies to both observational and action case studies. But an action case study should additionally describe exactly what the researcher did in the case organization, who were the stakeholders affected by the researcher's action, what their goals were, how the action was performed, what the outcome was and how this was evaluated by the stakeholders with respect to their own goals.

Finally, as any empirical research paper, a case study paper should contain a discussion of validity of its results, which includes an assessment of the quality and limits of the support for its observations and claims, and an acknowledgment of uncertainties that still exist about the results.

Observational case studies are typically done to increase our knowledge and understanding of a problem in practice, or to become aware of the variety of solutions to a problem actually used in practice. Action case studies are typically performed to test a new technique under conditions of practice. They have been variously called "pilot studies" (Glass 1997),

"technical action research" [5] or simply "case studies" [6]. However, they are very different from observational case studies because in an action case, the researcher intervenes in the case under study, whereas in an observational case, the researcher avoids disturbance of the case. Action cases provide the opportunity to learn from the confrontation of a technique with uncontrolled conditions of practice, but they have the danger of reduced generalizability, which should be properly acknowledged and discussed in the paper.

A case study report satisfying the above criteria would probably be rejected for an industrial conference. Even though relevance for research my be high, relevance for industrial practice might be very low or even absent. The case study may describe phenomena that the practitioner is already familiar with. The attention to research questions, conceptual framework, related work, measurement techniques, and validity discussions are pointless for the practitioner who wants to learn something new to achieve their goals.

V. FROM INDUSTRY TRACKS TO CASE STUDY TRACKS

Case study reports are a far cry from the partial and one-sided success stories that are too often submitted to industry tracks of research conferences. We should make room for industry case studies at research conferences if they satisfy stringent requirements of researcher papers such as the ones listed above. There is no reason to call them *industry* case study tracks, as in our field there is no other place than industry for case studies to be performed.

Separate case study tracks would have the benefit of making researchers aware of the complex conditions of practice for which their technical papers attempt to offer solutions. They would also serve a purpose in reporting about scaling up their technical solutions from the simplified and controllable conditions of the laboratory to complex and uncontrollable conditions of practice. A case study track would help making our solutions more relevant for practice in two ways:

1. Researchers can learn about practical problems and existing solutions from observational case studies published in such a track;
2. A case study track would help make our solutions more useful in practice and credible to practitioners, as it allows us to learn from pilot studies where solutions were tested under conditions of practice.

VI. RE@21: WHERE HAVE WE BEEN?

Tabulating the paper categories announced in the calls for papers in the ICRE and RE conferences since 1996, reveals the following trends concerning research and industrial tracks.

- The category of *Experience Paper* has been recognized since ICRE 1994, and is described in the calls for papers in terms that, with hindsight, indicate it as the category of an industrial experience paper. In RE 2002 (Essen), this category was called "industrial presentation" and since 2004 (Kyoto), the term "industrial practice and experience" became common. In 2011 (Trento), this was split into industry practice and industry problem subtracks.

These are the categories of paper of which we argue that it should be removed from a research conference like RE, and for which researcher's would be better off visiting industry conferences.

- Other categories occasionally mentioned under the industry track categories, are "ripe fruits" (ICRE 1994, 1996), "empirical results" (RE 2004), "Reflection on practice" (2005) and "applied research" (2013). From the descriptions of these topics in the CfPs we classify all of these approximately as *Technical Action Research,* a case study method in which the researcher uses a novel technique in practice to solve a real-world problem, or has practitioners use it to solve a problem, after which the researcher evaluates this instance of real-world use [5]. We think this should be included in the research track, and be evaluated against criteria for empirical research papers.
- In 2004, the industry track also contained a case study subtrack. The description in the CfP makes clear that *observational case studies* are intended, in which the researcher observes practice, analyzes the observations, explains them in terms of general theory, and then speculates about generalizability to other cases. This category too should be included in the research track.
- The category of *Research Evaluation paper* was introduced in RE 2004 (Kyoto) and has been stable ever since. From the start, Research Evaluation papers have included logical and formal analysis of techniques, so this category is not restricted to empirical research.

  The description of empirical research evaluation papers has varied over the years. It usually includes experiments, and at times the CfP also indicated simulations and case studies. However, surveys are not included, witness the separate category of survey papers that occasionally occurs in CfPs. From the CfP's it appears that survey papers should review the state of the art and provide a roadmap for research.

VII. ANALYSIS AND FINAL RECOMMENDATIONS

To sum up, the industry track contains some paper categories that we find important, viz. Technical Action Research and Observational Case Studies. Both categories are research papers in our view, next to empirical papers using experiment or simulation as a research method. The category of survey papers listed regularly in CfPs seems to mix three separate kinds of papers:

- Survey of the practice (using interviews, questionnaires etc. on a large sample of real-world cases).
- Meta-research (systematic review of the literature to aggregate research results) [7].
- Vision papers (roadmapping an aspect of RE research)

Only the first two can be classified as research papers. Vision papers are like opinion papers and have always been treated separately.

Combining this with our observation that scientific evaluation papers include both empirical and mathematical/logical

papers, this leads to three kinds of research papers:

- Technical solution papers
- Mathematical/formal research papers
- Empirical research papers
    - Experiments
    - Simulations
    - Observational case studies
    - Technical action case studies
    - Surveys of the practice
    - Meta-research
    - ...

The dots indicate that, although this list includes all empirical research methods mentioned in the RE and ICRE CfPs, we do not claim exhaustiveness.

With such a list of paper categories, an industry track is not needed. The list suggests an expansion of the conference to three kinds of tracks, each with their own evaluation criteria: technical solutions to be evaluated on novelty and relevance, mathematical papers to be evaluated on consistency, originality and added insight, and empirical papers to be evaluated on empirical research methodology and added insight. It goes without saying that a research conference is built on critical peer reviewing.

We claimed in the introduction that researchers who want to find out the state-of-practice in industry should attend industrial conferences. After all, if you want to know how Romans live, you visit Rome rather than invite Romans to visit you and give a couple of lectures ...

REFERENCES

[1] R. Yin, *Case Study research: Design and Methods*. Sage Publications, 2003, third Edition.

[2] P. Runeson and M. Höst, "Guidelines for conducting and reporting case study research in software engineering," *Empirical Software Engineering*, vol. 14, pp. 131–164, 2009.

[3] P. Runeson, M. Höst, A. Rainer, and B. Regnell, *Case Study Research in Software Engineering: Guidelines and Examples*. Wiley, 2012.

[4] R. Wieringa, N. Condori-Fernandez, M. Daneva, B. Mutschler, and O. Pastor, "Lessons learned from evaluating a checklist for reporting experimental and observational research," in *International Symposium on Empirical Software Engineering and Measurement (ESEM)*. IEEE Computer Society, september 2012, pp. 157–160.

[5] R. Wieringa and A. Morali, "Technical action research as a validation method in information systems design science," in *Seventh International Conference on Design Science Research in Information Systems and Technology (DESRIST)*, K. Peffers, M. Rothenberger, and B. Kuechler, Eds. Springer, 2012, pp. 220–238, lNCS 7286.

[6] B. Kitchenham, L. Pickard, and S. Pfleeger, "Case studies for method and tool evaluation," *IEEE Software*, vol. 12, no. 4, pp. 52–62, July 1995.

[7] B. Kitchenham, "Procedures for performing systematic reviews," Keele University/National ICT Australia, Tech. Rep. TR/SE-0401/0400011T.1, 2004.