Why it is Hard to Identify Technical Research Problems in ICT4D and How to Make it Easier

William Thies
Microsoft Research India
thies@microsoft.com

ABSTRACT

My position is that a shortage of detailed and compelling problem statements is the primary bottleneck that prevents most computer scientists from conducting research in ICT4D. While interesting problems exist, they are usually discovered via months of fieldwork, and there is little incentive to formalize and disseminate problems for the benefit of other researchers. To address this bottleneck, I argue that we should create a prestigious venue for publishing problem descriptions, rather than problem solutions. I also propose that we establish problem-exchange websites to solicit problems from practitioners; organize structured design contests that aggregate knowledge in a problem area; and leverage the domain knowledge of funding agencies in defining technical research problems.

1. INTRODUCTION

I believe that the primary barrier to the growth of ICT4D as an academic discipline within the field of computer science is that the burden on researchers is too high: they need to not only devise an innovative technical system, but also to understand a social problem to a depth that is not easily attainable. To lower the entry barrier to ICT4D research, I propose four mechanisms by which we can create and share detailed problem descriptions that computer scientists can use as a starting point for their research: 1) a publication venue that is devoted to problems rather than solutions, 2) a website that solicits problem descriptions from practitioners, 3) structured design contests that engage students and researchers in building knowledge around a problem, and 4) relationships with funding agencies that leverage their problem-specific knowledge.

My position stems from the observation that there is no shortage of computer scientists that are interested in working on technical problems that could impact global development. Almost everyone supports the mission and desires to get involved; however, they do not know what they could contribute within the realm of their technical expertise. Those that are skeptical of the research area are not (as a general rule) doubting the technical prowess of the researchers involved, but rather the technical depth and potential impact of the problems that they are working on.

2. WHY IT IS HARD

While formulating an interesting research problem is challenging in any field, it is unusually challenging in ICT4D because the problems addressed are not ones that are seen on a daily basis by researchers in high-income countries. In ICT4D, researchers are often trying to understand the problems that affect a very different population: one that lives in a different place, embraces a different culture, speaks a different language, and is subject to different economic, social, and technical constraints. While many problems in computer science are inspired by our daily experience in high-income countries – or at least are informed by colleagues elsewhere on campus (e.g., the department of biology) – the problems in ICT4D may never come up in our daily lives, or even in our newspapers or Internet browsers. More often, they are discovered and understood via experience on the ground.

Currently, researchers in ICT4D face two basic options for discovering problems of interest: they can visit the field themselves, or they can dialogue with a community partner. Of course, they can also follow their own intuitions regarding the relevance of their technologies; however, such strategies can lead (at best) to their own wasted efforts, as their solution is not relevant or adopted, and (at worst) to bad outcomes, as donors follow mis-guided enthusiasm to scale-up inappropriate solutions.

Researchers who have the flexibility and commitment to spend time in developing regions have the advantage of seeing the ground realities for themselves, and for relying on their own instincts and abilities for uncovering interesting research problems. However, there are also many drawbacks. It is costly and logistically difficult for most academics to make regular trips to developing regions. It is also difficult to expose the underlying realities during short visits, both due to time constraints and due to local bias (showing important visitors the best side of every coin). Moreover, the social skills required for effective ethnography and fieldwork often have little overlap with the quiet competencies required of a computer scientist. Inability to communicate in the local language may also pose significant barriers.

The second option is to rely on a community partner to relay information to the researcher (perhaps in addition to their field visits). This has the advantages of leveraging more experience than researchers could ever accumulate themselves; it also offers benefits beyond problem identification, such as providing feedback on prototypes and possible deployment of solutions. However, there are also drawbacks of relying heavily on a community partner. It is rare to find strong partners who understand the scope and potential of computer science research, and are interested in thinking on long-term time scales that are necessary for fundamental research. Researchers may also be at the mercy of the judgment and insights of the partner; if partners are imprecise in formulating the problem, then the research may also suffer. Finally, it may not be sustainable for researchers to collaborate with community partners indefinitely, as partners can potentially lose interest after investing time with researchers who did not have the capacity (or good luck) to provide a working solution to their problems.

To compound the concerns above, even when one does succeed in identifying a technical research problem in ICT4D, it is often outside one's own area of technical expertise. For example, if a networking expert spends months in the field, they may discover a research opportunity, but in speech technologies rather than networking. This phenomenon also represents an opportunity, as researchers that do not plan to pursue a discovered problem may be more willing to share that problem with colleagues. All that remains is to incentivize researchers to invest the effort needed to rigorously define and share the problems that they discover – a goal we address in the next section.

3. HOW TO MAKE IT EASIER

I propose four new approaches for fostering the identification and sharing of compelling research problems in ICT4D.

1. Reward problem statements with publication. I propose that we solicit detailed, contextualized, and unbiased statements of technical problems in developing regions for publication in conferences and workshops. These can be invited both as full papers in existing conferences (ICTD, NSDR, etc.) or perhaps as a new workshop or online journal. The crucial role of such a venue is to reward researchers for disseminating the insights gleaned from fieldwork (and conversations with partners), whether or not they have an innovative technical solution to match.

For example, I could envision a paper that defines the problem of medication adherence: what are the reasons for non-adherence, what are the constraints of delivering and consuming medication in rural areas, and what is the role of stigma, incentives, costs, families, politics, geography, and other factors as they relate to possible technical interventions to improve adherence. Such a detailed description of the problem is more than enough to fill a paper, but yet is often the minimum knowledge needed to undertake a technical research agenda in this area.

Researchers will reap rich returns from publishing such knowledge, as follow-up work on actual solutions is likely to cite their problem statement. If the authors have also developed solutions to the problem, then the solution is still eligible for publication elsewhere. If the problem is beyond their expertise, then the authors still get credit for defining it. Devoting an entire paper to the problem itself also encourages a rigorous an unbiased formulation, with less incentive to skew the description to match a particular solution. Other researchers – both inside and outside of ICT4D – would benefit greatly by having a single resource to consult for a breakdown of interesting problems in the field.

Publication of problem descriptions in ICT4D is analogous to publication of benchmark descriptions in computer architecture. Conferences such as the IEEE International Symposium on Workload Characterization are devoted to the subject.

2. Maintain a website of open problems. This website would provide a more informal and evolving portal to the same information contained in the publications above. However, as opposed to descriptions from researchers, the website would also solicit entries from practitioners, who often understand a problem deeply but are unable to cast it in terms that are interesting and appropriate for academic computer scientists. Such entries would be constructed with the assistance of a group of moderators, which would be drawn from respected researchers in the field.

As an example, consider a problem posed to us by Operation ASHA, a highly effective tuberculosis program that operates in New Delhi, India. They are seeking a means to reliably authenticate that a health worker and patient interacted at a given time. The solution must be low cost, as reliable as biometrics, and must offer timely notification (within a few hours) of each interaction. There are many additional parameters and constraints that deepens the problem definition. A range of solutions are possible, and are the subject of upcoming research.

In addition to the problem statement, the site would include updates regarding technical progress made on each challenge, as well as new demands and constraints from partners in the field. Such a resource would serve not only computer scientists who are looking for problems, but also for development agencies that are looking for solutions. A discussion board between problem solvers and solution seekers may also prove beneficial.

[1] `http://maven.smith.edu/~orourke/TOPP/`

This site would be analogous to the Open Problems Project[1] for computational geometry. It also bears some similarity to ThinkCycle.org (no longer available online), though with an emphasis on research problems rather than engineering design projects.

3. Organize structured design contests for students and researchers. I consider a *structured* design contest to be one in which all participants work on the same problem, and the organizers provide rich background materials that describe the context and constraints. In a classroom setting, a structured design contest has a place in any class that designs computing systems for the developing world. Such classes typically assign each student group to a different project, involving a large number of community partners and also burdening the students with finding a meaningful problem. However, there are many benefits to assigning the same problem to all of the students. The staff can invest deeply in building a knowledge base around the real constraints of the problem, including multiple perspectives from guest lecturers or from different community partners. Students can benefit by seeing others' approach to the problem. Also, partners benefit by choosing the best solution from the class, rather than bearing the risk of working with a single student team who might fail to deliver a working solution.

An example of a structured design contest is the Yunus Challenge to Alleviate Poverty, which is held at the Massachusetts Institute of Technology. Every year, students across campus can submit solutions to a single problem, in consultation with community partners, scientific experts, and other resources that are made available by the staff. Topics to date have included affordable small-scale energy storage systems, improving indoor cooking stoves, and improving adherence to TB medications. The competition has led to several spinoffs, including a non-profit co-founded by the author (Innovators In Health).

Structured design contests also have a role in the research community. Examples in other fields include the Multimedia Grand Challenges[2], the Supercomputing Challenge[3], and the CHI Student Design Contest. I think the time is right for a design contest in the ICT4D space, with a dedicated track at a premier conference.

4. Leverage the expertise of funding agencies. In other areas of computer science, researchers leverage the applications expertise of funding agencies to ground and direct their research (e.g., the DARPA Grand Challenge). The potential benefits are perhaps higher in the ICT4D domain, as global foundations have rich resources and knowledge bases in areas of international development that can be used to identify interesting research problems.

An example of this direction comes in the form of open, structured design contests that have recently been organized by leading foundations. Last year, the Rockefeller Foundation partnered with Innocentive to offer a $40,000 award to the best design for a solar power device that reduces the risk of malaria. While funding at this level may not sustain an academic research program, the foundation also provided open guidance in the form of literature reviews and other pointers which could direct ongoing research in the area. Other challenges on the site include "improving banking processes in the developing world" and "solar-powered wireless routers".

Engaging funding agencies with respect to specific research challenges could not only help to identify relevant problems, but perhaps also increase the chances of sustained funding.

Acknowledgments. I thank Saman Amarasinghe, David Hutchful, Indrani Medhi, Thomas Smyth, Kentaro Toyama, and Julie Weber for helpful conversations.

[2] `http://www.scils.rutgers.edu/conferences/mmchallenge/`

[3] `http://sc09.supercomputing.org/?pg=challenges.html`