Why Do We Need This Workshop?

James McCauley, Aurojit Panda, Scott Shenker

The purpose of this workshop is to bring together researchers who are interested in "designing computer systems and architectures in a socially responsible way" [4]. This paper starts with simple question: Why do we – as researchers who are interested in such issues – feel the need for a special workshop devoted to this topic? To wit, why isn't social responsibility important in all technical conferences, and treated as just another metric along which systems are evaluated? Our response to this question consists of two obvious observations, that in turn lead us to the more fundamental question: How can this emerging "socially responsible" community have impact?

Our first obvious observation is to note that the phrase "socially responsible" typically refers to system properties that are beneficial to society but perhaps not as much to the individual companies that deploy these systems. We don't need special workshops on how to build systems that are faster or more cost-efficient, because these are properties that provide direct benefit to the deploying companies. However, various other properties – ranging from bias and fairness to carbon footprint and lack of accessibility – give rise to negative (or positive) externalities that are felt more by society than by the deploying companies, and these factors are often not taken into account when commercial deployment decisions are made.

But this does not explain why we need a special workshop because we, as a research community, could still focus on these socially-relevant properties even if they are of less interest commercially. So the uncomfortable question we must face is: why aren't we focusing on this work in our traditional conferences?

Our second obvious observation is that the work we do as a research community is strongly influenced by what is of interest to commercial companies. This is not due to nefarious reasons (such as having financial conflicts of interest), but arises quite naturally from our desire to have impact. Many researchers, particularly those interested in systems design, want to do work that will influence what eventually gets deployed. Commercial companies are typically the ones making the decisions about what systems are deployed on a large-scale basis, so their influence on our research agenda is understandable. Of course, there are many other factors that influence the work we do, including what conferences will accept, and what government and industry will fund, but these tend to reinforce the emphasis on impact – and therefore increase the influence of commercial companies – rather than mitigate it.

So what are the implications of these two observations for socially responsible research and the goals of this workshop? For a researcher interested purely in deployment of their ideas, they might evaluate their work in terms of Probability of Impact on Deployment (PID), which measures the impact the research would likely have on deployed systems; it captures whether the paper changed the way people build deployed systems, regardless of whether the specific system described in the paper is ever deployed. Using this as a criterion for choosing work would give tremendous power to commercial companies to set the research community's agenda, because they are the gatekeepers for most large-scale deployments. A cynical view of our current conferences might suggest that this is the prevailing metric in many communities.

If we want to create a community of socially responsible researchers, what prevailing metric for evaluating work should we adopt? One might think the natural metric would be to evaluate work in terms of its Benefit to Society (BtS). This may lead to a literature full of interesting and beneficial ideas, but with little chance of adoption because none of the factors that drive commercial adoption decisions have been taken into account.

Another possibility is that researchers evaluate their work by the product of PID and BtS; that is, the probability the work will impact deployments times the societal benefits such impact would bring. While this formulation does include societal considerations, it still allows the commercial gatekeepers to have a strong influence on the research agenda because for many socially beneficial technologies the first term would be close to zero.

Thus, with these two options we have unappealing choices, either doing work without regard to its deployability or giving veto power to the commercial gatekeepers. We contend that for such a socially-responsible research community to have a beneficial impact on society, it should not accept deployability as exogeneously decided but instead seek to improve the chances of adoption; in short, it should view adoption as another engineering

design exercise.

More specifically, we urge this community to embrace a research agenda that has two tracks. The first track is similar to the call for this workshop, focusing on the development of new technologies that are superior to what we have today, but where the superiority is evaluated in terms of socially-relevant criteria. We, as technology researchers, know how to pursue this kind of technically-motivated research agenda, and having a conference where this is the criterion for publication would be valuable.

However, to have an impact, we should embrace a second track of research that involves exploring how to change technology ecosystems so that companies might be more likely to adopt socially beneficial technologies. In a previous paper [2] and expanded on in a later talk [3], we have described how we might address ecosystem problems through an approach called Technology Ecosystem Transformation (TET). TET works as follows: if one wants to deploy some particular socially-beneficial technology, one first tries to identify a technical intervention with the following two properties:
(1) It can be deployed without requiring the cooperation of the dominant incumbents. The deployers may be smaller commercial companies, or nonprofits, or other organizations that have an incentive to deploy this intervention for its own sake. Finding the deployers is often as difficult as identifying the intervention.
(2) The presence of this technical intervention in the ecosystem changes the incentives for incumbents, so that they are now motivated to deploy this particular socially-beneficial technology.

With such an intervention, one first deploys the intervention and then lets the resulting market incentives drive the adoption of the socially responsible technology. This general approach is well-known when the intervention is some form of regulation or tax, such as when governments invoke penalties or provide rebates to reduce pollution or encourage recycling. The TET approach uses technical interventions to provide the appropriate positive or negative incentives. The TET approach can coexist with other forms of incentives, but does not require them. We give two recent examples of the TET approach.

The first, described in [2], is an attempt to provide some degree of privacy for personal photos by preventing their spread on social media sites if they are maliciously or mistakenly posted. The intervention consists of (i) a backend "ledger" of photos, (ii) device software that signs every photo and labels them as sharable or not, and (iii) browser mechanisms to check whether a photo is sharable. This intervention can be deployed by the entities providing privacy-oriented browsers (e.g., Brave, Mozilla) for relatively little cost, and yet scale to billions of photos. At that point, the technology is well-enough entrenched that incumbents would have an incentive to adopt because photos marked as not-shareable could be trivially recognized by any major website, and thus non-adopting sites could be liable for any damages that result from such sharing. In addition, some incumbents might want to make privacy a selling point by explicitly adopting this approach. Lastly, the presence of such technology would allow regulators to start requiring its use. So the intervention simultaneously creates positive incentives for adoption, negative incentives for not adopting, and a technical solution that regulators can require.

Another example is Sky Computing [1, 5], whose goal is to create a more competitive cloud market by mitigating against lock-in business practices. The intervention is the introduction of intercloud brokers that help create a two-sided market for cloud compute jobs. These intercloud brokers do not require the cooperation of clouds, and could charge a small fee for finding the cheapest or highest-performing cloud. We have already seen widespread interest in Sky Computing from the non-dominant clouds. When a significant portion of compute workloads are being handled by these intercloud brokers, clouds will have an incentive to make their offerings more compatible and focus less on lock-in business practices.

The TET approach requires, for a given socially desirable technology, finding an intervention that is both deployable and creates the right incentives for the adoption of the desirable technology. It has to be the "right" intervention, so this is a hard design problem, but it need not be a technically "clever" or novel intervention. That is not the kind of work our community has often pursued; as Don Norman said, "Academics get paid for being clever, not for being right." But as a community devoted to social responsibility, we should remember that being right is far harder, and far more important, than being clever.

Thus we, as a community of researchers interested in socially-responsible design, should pursue both *technical* problems (where we need to invent new socially beneficial technologies) and *deployment* problems (where we need to ease the deployment of socially beneficial technologies, whether or not achieving this deployment requires mechanisms that are clever or novel).

References

[1] Sarah Chasins, Alvin Cheung, Natacha Crooks, Ali Ghodsi, Ken Goldberg, Joseph E. Gonzalez, Joseph M. Hellerstein, Michael I. Jordan, Anthony D. Joseph, Michael W. Mahoney, Aditya Parameswaran, David Patterson, Raluca Ada Popa, Koushik Sen, Scott Shenker, Dawn Song, and Ion Stoica. The sky above the clouds, 2022.

[2] Narek Galstyan, James McCauley, Hany Farid, Sylvia Ratnasamy, and Scott Shenker. Global content revocation on the internet: a case study in technology ecosystem transformation. In *Proceedings of the 21st ACM Workshop on Hot Topics in Networks*, HotNets '22, page 115–121, New York, NY, USA, 2022. Association for Computing Machinery.

[3] Scott Shenker. Raising our sights (a long rant from an accidental engineer). https://www.youtube.com/watch?v=XLZ0sVbA3cs, December 2023.

[4] Akshitha Sriraman. Invitation to submit to the first hotethics workshop (co-located with asplos '24). email inviting submissions, March 2024.

[5] Ion Stoica and Scott Shenker. From cloud computing to sky computing. In *Proceedings of the Workshop on Hot Topics in Operating Systems*, HotOS '21, page 26–32, New York, NY, USA, 2021. Association for Computing Machinery.