

Editor's Message: The Goal of Systems Research

What is, or ought to be, the goal of systems research? The answer to this question differs for academics and researchers in industry. Researchers in the industry usually work either directly or indirectly on a specific commercial project, and are therefore constrained to design and build a system that fits manifest needs. They do not need to worry about a goal beyond this somewhat narrow horizon. For instance, a researcher at Google may be given the task of building an efficient file system: higher level goals beyond this are meaningless to him or her. So, the 'goal' of systems research is more or less trivial in the industrial context.

Many academic researchers in the area, however, are less constrained. Lacking an immediate project to work on, they are often left wondering what set of issues to address.

One solution is to work with industrial partners to find relevant problems. However, although this results in problems that are well-defined, immediately applicable, and even publishable in the best conferences, it is not clear whether this is the true role of academia. Why should industrial research be carried out for free by academics, in effect subsidized by society? I think that academics may be inspired' by industrial problems, but should set their sights higher.

Another easy path is to choose to work in a 'hot' area, as defined by the leaders in the community, or a funding agency (more often than not, these are identical). If DARPA declares technology X or Y to be its latest funding goal, it is not too hard to change ones path to be a researcher of flavour X or Y. This path has the attraction that it guarantees a certain level of funding as well as a community of fellow researchers. However, letting others decide the research program does not sound too appealing. It is not that far from industrial research, except that the person to be satisfied is a program manager or funding agency, instead of your boss.

I think academic researchers ought to seek their own path relatively unfettered by considerations of industrial projects or the whims of funding agencies. This, therefore, immediately brings up the question of what ought to be the goal of their work. Here are my thoughts.

I believe that systems research lies in bridging two 'gaps': the Problem Selection Gap and the Infrastructure-Device Gap. In a nutshell, the goal of systems research is to satisfy application requirements, as defined by the Problem Selection Gap, by putting together infrastructure from underlying devices, by

solving the Infrastructure-Device Gap. Let me explain this next.

What is the Infrastructure-device gap? Systems research results in the creation of systems infrastructure. By infrastructure, I mean a system that is widely used and that serves to improve the daily lives of its users in some way. Think of it as the analogues of water and electricity. By that token, Automatic Teller Machines, Internet Search, airline reservation systems, and satellite remote sensing services are all instances of essential technological infrastructure.

Infrastructure is built by putting together devices. By devices, I actually mean sub-systems whose behaviour can be well-enough encapsulated to form building blocks for the next level of abstraction and complexity. For instance, from the perspective of a computer network researcher, a host is a single device. Yet, a host is a complex system in itself, with many hundreds of sub-systems. So, the definition of device depends on the specific abstraction being considered, and I will take it to be self-evident, for the purpose of this discussion, what a device is.

An essential aspect of the composition of devices into infrastructure is that the infrastructure has properties that individual devices do not. Consider a RAID system, that provides fault tolerance properties far superior to that of an individual disk. The systems research here is to mask the problems of individual devices, that is, to compose the devices into a harmonious whole, whose group properties, such as functionality, reliability, availability, efficiency, scalability, flexibility etc. are superior to that of each device. This then, is at the heart of systems research: how to take devices, appropriately defined, and compose them to create emergent properties in an infrastructure. We judge the quality of the infrastructure by the level to which it meets its stated goals. Moreover, we can use a standard 'bag of tricks' (explained in the networks context in George Varghese's superb book 'Network Algorithmics') to effect this composition.

Although satisfying, this definition of systems research leaves an important problem unresolved: how should one define the set of infrastructure properties in the first place. After all, for each set of desired properties, one can come up with a system design that best matches it. Are we to be resigned to a set of not just incompatible, but incomparable, system designs?

Here is where the Problem selection gap fits in. Systems are not built in a vacuum. They exist in a social

context. In other words, systems are built for some purpose. In the context of industrial research, the purpose is the purpose of the corporation, and handed down to the researcher: ‘Thou Shalt Build a File System’, for instance. And along with this edict comes a statement of the performance, efficiency, and ‘ility’ goals for the system. In such situations, there is no choice of problem selection.

But what of the academic researcher? What are the characteristics of the infrastructure that the academic should seek to build? I believe that the answer is to look to the social context of academia. Universities are supported by the public at large in order to provide a venue for the solution of problems that afflict society at large. These

are problems of health care, education, poverty, global warming, pollution, inner-city crime, and so on. As academics, it behooves us to do our bit to help society solve these problems. Therefore, I claim that as academics, we should choose one or more of these big problems, and then think of what type of system infrastructure can we build to either alleviate or solve it. This will naturally lead to a set of infrastructure requirements. In other words, there is no need to invent artificial problems to work on! There are enough real-world problems already. We only need to open our eyes.

S. Keshav
CCR Editor

