So Spencer gave some good, concrete examples of how software engineering research can have an immediate impact on climate science and data sharing.

Now, I want to broaden the topic. To frame a discussion about what we as software engineering researchers can do, we first need to look at what the challenge problems are, and what intellectual assets we bring to the table. What are we good at? What do we know how to do?

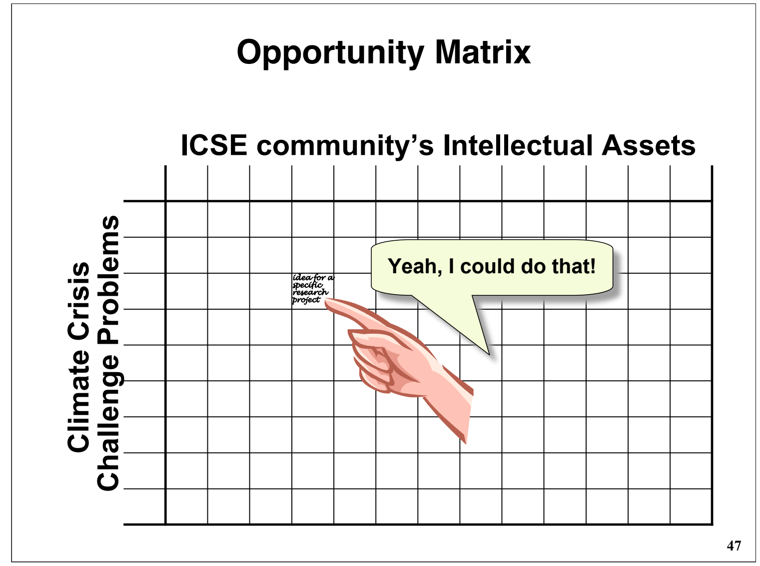

Here's one way to look at it. We identify a list of major challenges: systemic changes that have to happen; new technologies that have to be developed and deployed, and so on. Then we look at all the things that software engineering, as a discipline, has to offer. Then, we attempt to populate the matrix, identifying specific research projects that people in the ICSE community will get excited about as research projects.

So, what do I mean by our "intellectual assets"? Let me suggest a few…

Asset: Abstraction, Separation of Concerns

image credit: lumakick.com

48

We have analysis and modeling techniques for thinking about complex designs, and building useful abstractions. Note: this is something that many scientists (especially physicists) are lousy at. They focus on detail, using the reductionist approach that dominates the natural sciences. They are poor at systems thinking, and poor at building whole system conceptualizations and visualizations. Abstraction is so central to the way software engineers think that we often don't realise that it's one of our major strengths as a discipline, and other disciplines don't share it.

Asset: Modularity

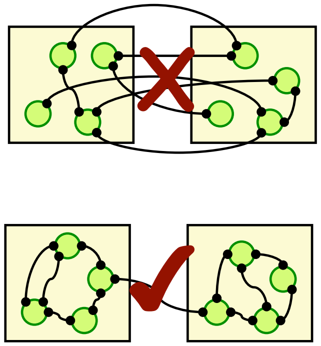

49

We know how to design information hiding modules to create architectures for managing change in very complex systems. Again, this is core to our discipline, and again it's something that other disciplines know very little about. Our ability to design systems that minimize coupling between modules, to support future system evolution will be an essential contribution for many of the challenges.

Asset: Open Source Culture & Practices

image credit: Today is a good day @flickr

We know about very successful practices and tools for large scale open collaboration on complex socio-cognitive design tasks. We know a lot about the sociology of these communities, and the way in which they operate as meritocracies. We know a lot about the “informalisms” that these communities use to create enough shared understanding of a design, but without having to meet face to face, nor create extensive design documents. We also know a lot about how these communities remain agile.

Asset: Requirements Modelling

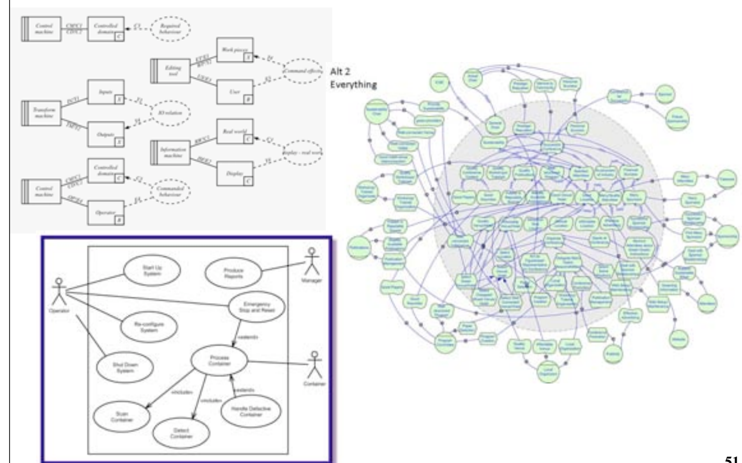

We know how to analyze complex multi-stakeholder socio-technical problems, identify goals, decompose problems, etc. We have techniques for eliciting and negotiating requirements from diverse stakeholders, and for modeling and validating these requirements. Many of the challenges depend on ultra-large scale information systems. Advanced systems analysis techniques will be needed to help us understand these systems from multiple perspectives.

Asset: Formal Verification Techniques

$$
\begin{array}{|l}
project : OZSpec \rightarrow UMLDiagram \\
\hline
\forall (oz, uml) : project \bullet \\
\quad \{c : oz \cap Classdef \bullet c.name\} = \{c : uml.classes \\
\qquad \bullet c.name\} \bullet \forall c_1, c_2 : oz \cap Classdef \bullet \exists_1 c' : \\
\qquad\quad uml.classes \bullet c'.name = c_1.name \\
\qquad\quad c'.attris = \{cls : Classdef \mid cls \in oz \bullet cls.name\} \\
\qquad\qquad \lhd c_1.state.decpart \\
\qquad\quad c'.ops = \{o : Opdef \mid o \in c_1.ops \bullet o.name\} \\
\qquad c_2.name \in \{t : \mathrm{ran}\ c_1.state.decpart \bullet t.name\} \Rightarrow \\
\qquad\quad \exists_1 (c_1', c_2') : uml.agg \bullet c_1'.name = c_1.name \\
\qquad\qquad \wedge c_2'.name = c_2.name \\
\qquad c_2.name \in \{inh : \mathrm{dom}\ c_1.inherit \bullet inh.name\} \Rightarrow \\
\qquad\quad \exists_1 (c_1', c_2') : uml.inh \bullet c_1'.name = c_1.name \\
\qquad\qquad \wedge c_2'.name = c_2.name
\end{array}
$$

52

We know how to describe, design and verify real time mission-critical control systems - as we move towards smart control over energy production and consumptions, we'll need even more sophisticated control systems. These will tax our ability to build reliable, predictable real-time software.

Asset: Distributed Collaborative Design

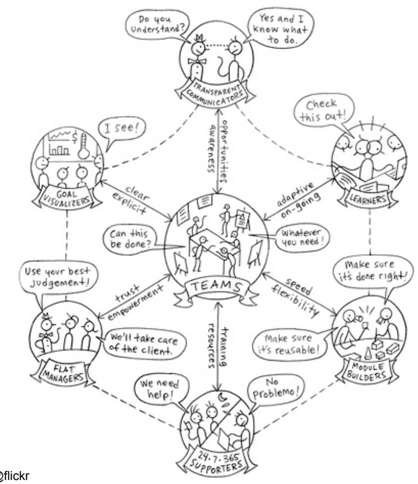

image credit: dgray_xplane @flickr

53

We know how to match team and organizational structures with technical system architectures, to support collaboration, especially across geographic distances. For example, the work of the Socio-Technical Congruence (STC) workshop at ICSE, looks at Conway's law, and the use of social network theory to provide awareness tools, visualizers and recommenders, to make sure the right people talk to each other to coordinate their contributions to the project.

Asset: Mining Repositories

image credit: found drama @flickr

We know a lot about data mining and visualization, especially for time varying data on the evolution of technical and social structures. For example, the work of the Mining Software Repositories (MSR) workshop, in which many sophisticated tools are being developed, but for which convincing applications are often still missing.

Asset: Self-Healing Systems

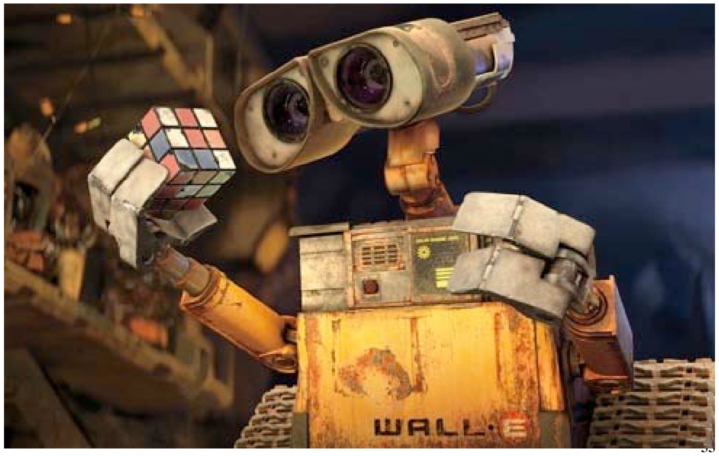

We have techniques for monitoring, repairing and upgrading operational systems, in complex environments where the system cannot be taken down for repair. We know how to build systems that can diagnose their own failures, and take automated repair actions. And, more prosaically, we know how to incrementally improve operational systems via upgrades and patches.

Asset: Debugging Strategies

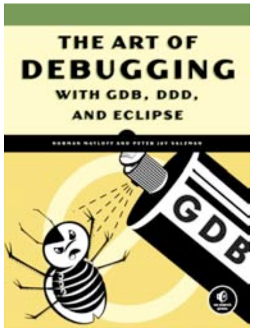

56

We know a lot about systematic strategies for tracking down errors in complex designs. We're good at spotting feature interactions, characterizing programmer errors, boundary errors, etc. This may be a part of what Jeanette Wing calls "Computational Thinking". We techies are brilliant at spotting and diagnosing problems in human organisational systems that other people never notice. Many of our debugging strategies could be applied to debugging other types of system: e.g. social systems, policy frameworks, etc.

Asset: (the idea of) Design Patterns

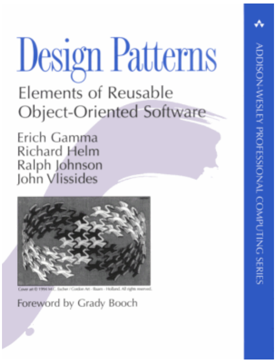

57

We're good at identifying design pattern as (1) *proven solutions* to a (2) *commonly occurring problems* in (3) *particular contexts*. The idea of identifying patterns and creating patterns books didn't originate in this community, but we've taken it to a whole new level.

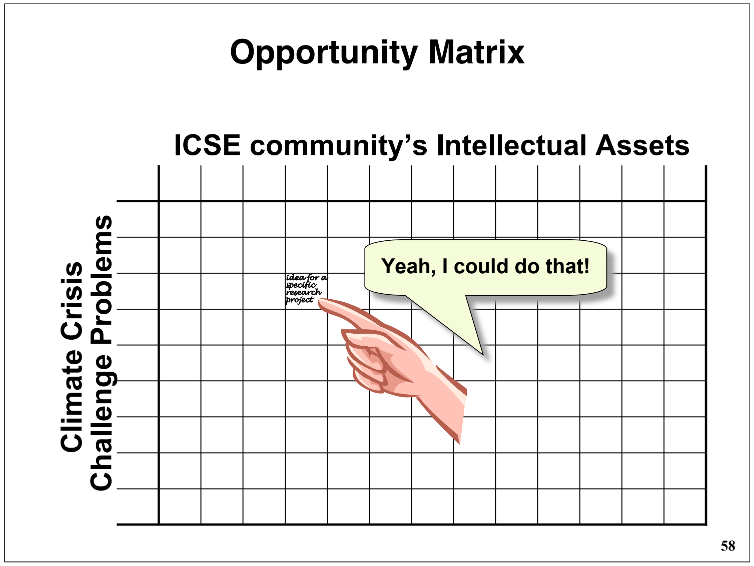

Okay, so that's a few samples of what I mean by intellectual assets. There are many others.

What about the challenge problems? One way to think of them is in terms of the key audience for research we might do: climate scientists, policymakers, journalists, educators, engineers, etc. Let's take a look at a few examples…

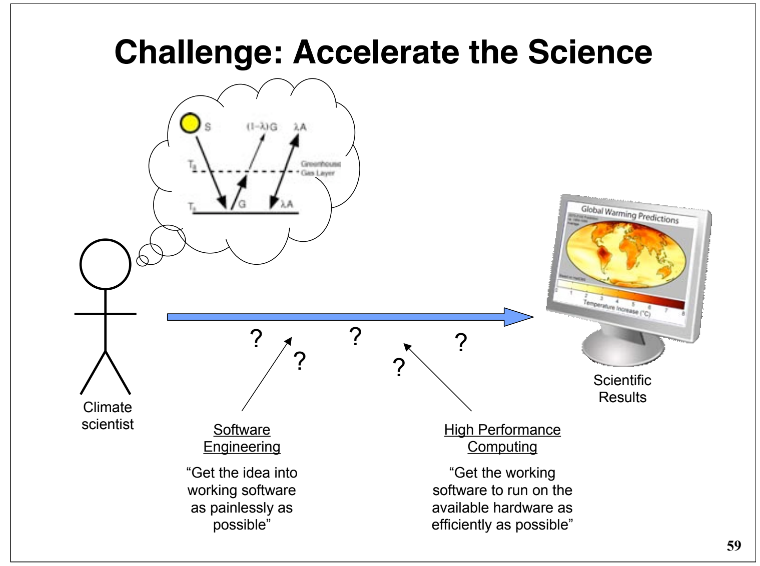

First, the climate scientists themselves. Much research has gone into effective use of high performance computing, which is one of the critical bottlenecks in developing complex simulations of earth systems. But the other bottleneck, software development for scientific computation, has barely been studied at all. Improving the software practices of these scientists might offer large productivity gains.

Challenge: Scientific Data Sharing

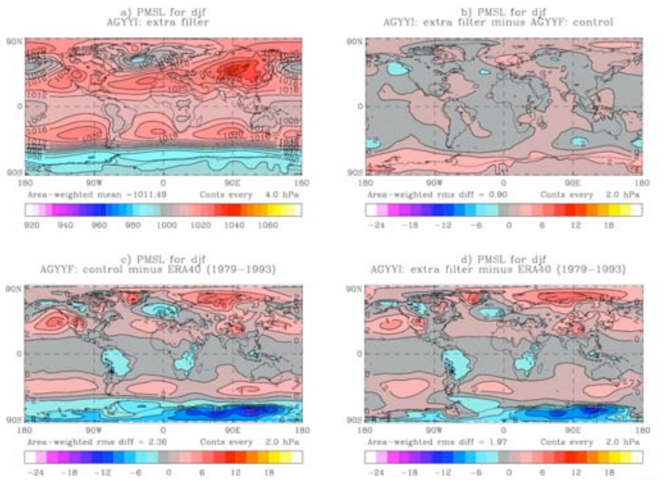

60

Climate models and other earth system models create huge volumes of data, which is hard to understand and hard to use. Yet there is a rapidly increasing set of other scientific communities who are potential downstream users for this data. These communities lack the meta-data standards by which they can find out what data exists, and what assumptions are embodied in those data sets. The need to develop shared ontologies across disciplinary boundaries, shared domain models, use cases to capture their requirements, and the technical infrastructure to make all this work.

Challenge: Open Collaborative Science

61

Beyond just the data, we lack the tools and cultures around open science, especially for collaborating across disciplinary boundaries. For example, how do scientists identify relevant work related to their own as it is published? How do they share the tacit knowledge surrounding their data, procedures and results? How do we coordinate massive distributed data collection efforts? How should scientists use their social networks more effectively for knowledge sharing?

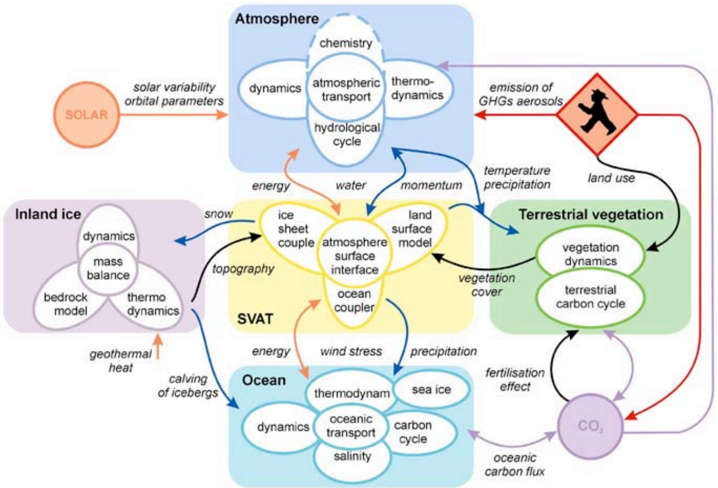

62

Just to elaborate on that point further: the current state of the art in earth system modeling incorporates simulation code from many different disciplines, from atmospheric science, oceanography, biology and the carbon cycle, chemistry, hydrology (the study of water cycles and ice dynamics), space sciences (solar radiation and top of the atmosphere effects) and human social systems. All of these subsystems are coupled together to study systemic effects in a whole system model. This image is a conceptual architecture for the current generation of earth system models.

Challenge: The Smart Energy Grid

63

We need to completely replace current electricity transmissions systems with a transactional power grid that's more like the internet. Energy will be produced at many micro-generation sites, and sold back to the grid at varying price rates. When everyone has solar panels and mini wind turbines, there will no longer be a distinction between producers and consumers. The smart grid will dynamically respond to shifting patterns of generation and production, and automatically adjust power usage to match supply in the many devices that can choose when to draw power and when to idle.

Challenge: The Electric Car

64

We need to build zero emissions vehicles and solve the complex software control problems that make electric cars as driveable as internal combustion cars, even in the face of variable power draw on the batteries, and variable supply over the charge cycle.

Challenge: Zero Carbon Buildings

Build zero emissions houses/condos, using a mix of passive energy (e.g. appropriate windows and solar walls) and active system that adjust the building to external conditions. These building will be filled with smart sensors that monitor energy production and consumption, and provide feedback to the users, and automated adjustments to building controls.

Challenge: Power Aware Computing

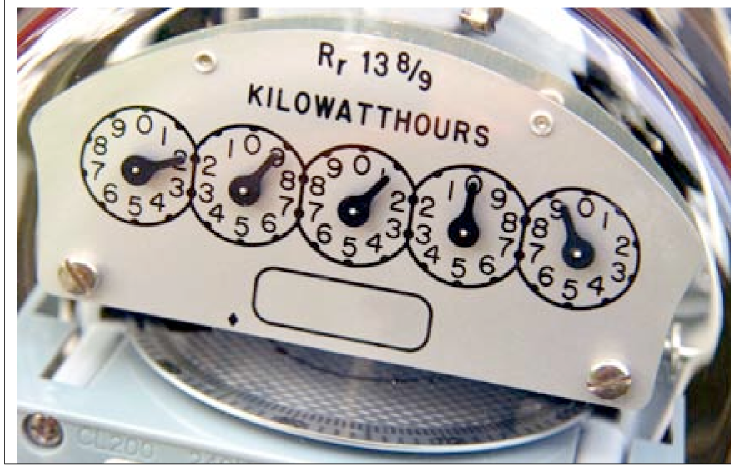

We need to build software applications that understand and adjust their own power consumption. We need to take the startling progress that has been made in power management for mobile devices and apply it to all computing. And we need to match the progress in energy efficient hardware with similar advances in energy efficient algorithms and systems software.

Challenge: Decision Support for Policymakers

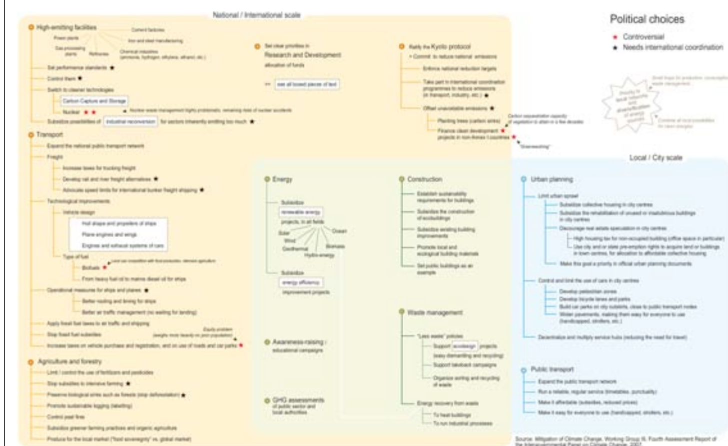

Here's another set of stakeholders: policy makers who need to make decisions and defend policy choices by linking them to the latest available science - a kind of evidence-based policy making. They need decsion support tools that indicate which of the various control levers are likely to be most effective for their defined targets. And they need to know how to adjust these policy tools once implemented, to incrementally adjust their effectiveness. Agile policymaking, anyone?

Challenge: Architecting Climate Policy

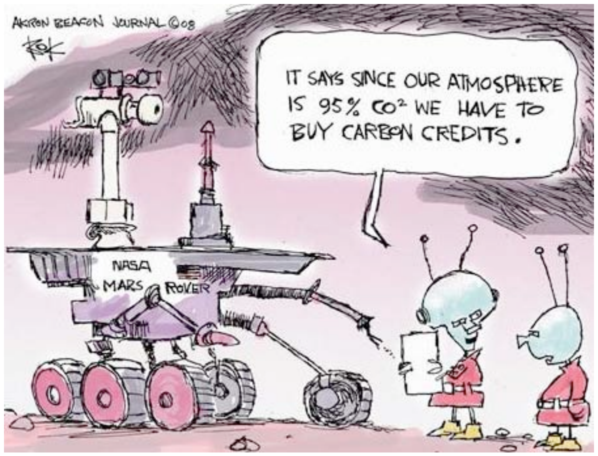

68

Here's an interesting challenge: How do we develop an appropriate architecture for the global climate policy, so that it has the desired effects on systems it is supposed to control, while minimizing undesirable effects on other systems (such as the health and well being of the people). Can we design a modular architecture for global policymaking using information hiding and loosely coupled modules? Can we identify good design patterns for such policy making?

Challenge: Validating Planning Models

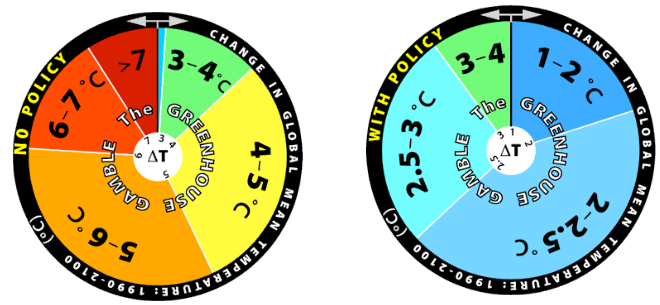

69

And we need to validate the models (economic, social, physics, etc) that inform policymaking. We need comprehensive, systematic verification and validation processes for these models, to ensure the evidence used in policymaking is sound. This image is an analysis using the MIT model, of the effects of a particular climate policy in shifting the probabilistic temperature forecast down. Is the underlying economics model valid? How would we tell?

Challenge: Amplify Grassroots Action

And one last group of stakeholders we can help. We need to help the advocates for change to push their governments into rapid, radical change. We can provide tools to facilitate social networking, design patterns for effective protest, open collaboration for activists, and so on. We can harness the power of the open source movement to make these activist movements more effective.

Also, a related item: we can raise the level of discourse among journalists and bloggers, by providing appropriate information tools in a form that matches their needs.

Opportunity Matrix

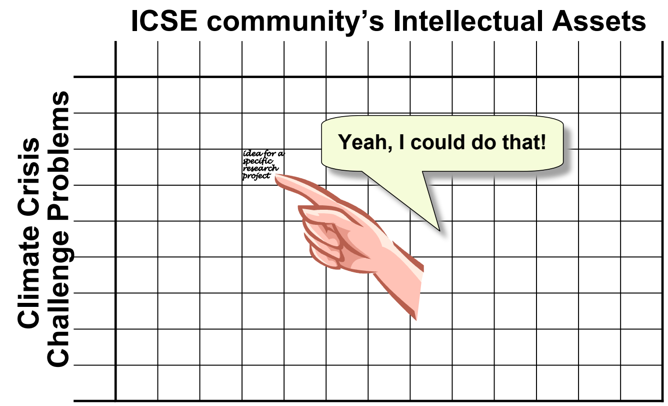

So, that was whirlwind tour of some of what I think of as the software engineering research community's intellectual assets, and a shortlist of some sample climate challenge problems. I'd like to propose we continue to develop a research agenda by exploring this matrix, and identifying the opportunities for the most effective action.

Discussion:
Build a Research Agenda

72

With that, I'd like to open up the discussion from the floor…