Distributed Group Problem Solving

Bruce Skarin
Senior Simulation Scientist, Aptima Inc.
bskarin@aptima.com

Abstract

There are a multitude of "grand challenges" facing the world, ranging from global climate change and energy needs to growing deficits and health care costs. The increasing urgency of these challenges brings into question the abilities of society to solve such complex problems. At the same time, the Internet and social media have come to provide unprecedented access to the millions of stakeholders. While this capability has helped facilitate changes as far reaching as the 2011 revolution in Egypt it is still unclear if it can be used as part of a broader constructive process. For grand challenges the process will need to be massively collaborative and yet well structured. This paper explores how system dynamics, combined with web technologies can provide both the workflow and framework required. A model-based approach is ideal because theories can be made more transparent and easily compared to real world data. To explore the concept of using the system dynamics methodology as part of a distributed group problem solving approach this paper also describes a work in progress demonstration web site. The site uses existing web applications to illustrate the necessary components for creating a problem solving process on a grand scale.

Introduction

Over the last few centuries, tremendous advances have been made in science and society. In developed countries chronic problems such as starvation and disease have been mitigated, governments have stabilized, equality and justice largely maintained, and education provided for. Yet in the march of this progress new, more complex problems have come into focus such as threats to security, economic vulnerabilities, and the potential for significant global climate change. These examples are part of a new class of problems that have been called "grand challenges." This is an appropriate name since they are both large in size and complexity. Such dimensions have challenged our conventional approaches to problem solving. Where policy development processes used to be effective, we now find deadlocked legislatures and polarized populations. Where there used to be time to experiment and discover, we now find mounting pressure to perform.

One relevant and active example is the latest effort to reform the United States health care system. In the past few years the debate over health care in the U.S. has identified numerous issues with the current system including: unchecked growth in costs, uninsured citizens, insured patients being denied care, and a declining cost to quality ratio. These issues have sparked high-quality discussions as well as emotional exchanges. They have incensed partisan politics even when there is bipartisan agreement on many of the core issues. As the debate continues, the expectations for reaching a consensus have diminished and the threat of a future impasse looms ever larger. This growing deficiency in complex problem solving is also not limited to governments, as demonstrated by the recent instabilities seen in financial, automobile, and energy industries following 2009. The question then is: if current approaches to solving problems are insufficient, what is needed to help create viable solutions?

Current barriers for problem solving

Approaches to grand challenges are plagued by two main issues: 1) their complexity, which is created by their sheer size and entanglement with large numbers of people and 2) their dynamics, which generate behaviors beyond individual comprehension. There are numerous examples of large, complex, and dynamic problems being managed successfully, such as putting a man on the moon. Yet these problems have the benefit of primarily dealing with physical systems. In dealing with a system like health care, reliable physics are scarce, and rules and procedures are subject to wide interpretation. Instead, problems with health care emerge from the actions of hundreds of millions of individuals whose behaviors are not always rational or consistent. While numeric and scientific analyses are certainly still necessary to develop better policy, as a purely engineering problem, a system like the U.S. health care is intractable. This is because, despite some of the more concrete economic factors, it is also invariably a social problem.

In working through social problems the federal constitutional republic system of the United States has successfully navigated numerous complex issues such as social equality, education, and justice. While struggles with these issues are by no means as complete and as smooth as engineering problems, the processes and frameworks in place ultimately have been able to deliver working solutions to these and other social problems.

Solutions to grand challenges will therefore need to be developed using a process that is both scientific and social. What is needed is system that allows for the rigorous treatment of a given challenge while allowing millions of people to participate in a manner similar to past social movements. Such an ambitious system will face three key challenges:

- Establishing a workflow that allows distributed contributors to decompose and analyze problems collaboratively
- Locating, motivating, and retaining a massive and wide ranging group of users from experts to the uninformed
- Preventing the manipulation of the system while allowing for diverse “outside of the box” thinking during the development of possible solutions

Emerging capabilities

When considering the engineering and social challenges of creating such a system, two promising developments have emerged over the last decade of internet fueled evolution. The first is the creation of new architectures for conducting distributed collaborative development and computation. Included in these creations are Wikis, open source code repositories such as GitHub, and distributed computing frameworks like BOINC (used by projects like SETI@home). The second development is the explosion in social networking sites and services. Examples include sites such as Facebook and LinkedIn and services such as Reddit and StumbleUpon. These two developments provide some of the essential building blocks for the required system in the form of technologies and frameworks for organizing and executing complex engineering tasks along with massive social networks comprised of millions of users.

Some visionaries have characterized this social and engineering transformation as the emergence of a global mind or consciousness. From this perspective the cognition of the global mind is then a distributed process where individuals create, analyze, transform, and redistribute information through a variety of Internet media. Currently, the majority of this global cognition might be considered as “shallow thinking” where the problems may be socially interesting but as inconsequential as the selection of the next American Idol. Yet there are also some instances of smaller scale “deep thinking” where journalists and bloggers who are also experts in a field seek to elucidate some of the intricacies of an issue like health care reform. The design challenge presented then is in how to bridge the gap between the larger shallow thinking networks and the smaller deep thinking networks.

This paper will describe a design concept that leverages data, models, and the system dynamics problem solving process to provide one possible framework for approaching grand challenges. To illustrate how the framework can be used to examine a problem, a work in progress study of health care reform in the United States will be used. This study will show how existing web applications can be combined to provide a majority of the functionality needed using a functioning demonstration site: www.socialsolver.com/project-definition/health-care. As a proof of concept, the site has remained closed as of this writing. As such, only the technology is discussed in detail, leaving many of the questions about the group collaboration unexplored. The conclusion will discuss some of the remaining technology gaps and next steps.

System design

The design of a system for distributed group problem solving can draw on the functionality of many existing web applications and services as shown in Figure 1. The majority of the system is encapsulated in a collaboration site that utilizes Wiki-style pages to organize content along a workflow that is modeled after the system dynamics problem solving process. Within the site other services are embedded to provide interactive descriptions of each project, data views, simulation models, as well as discussion and polling forums. By embedding other services like YouTube, Google Docs, and Facebook, the site will not only leverage existing functionality, it will facilitate the remixing of content with other sites and services.

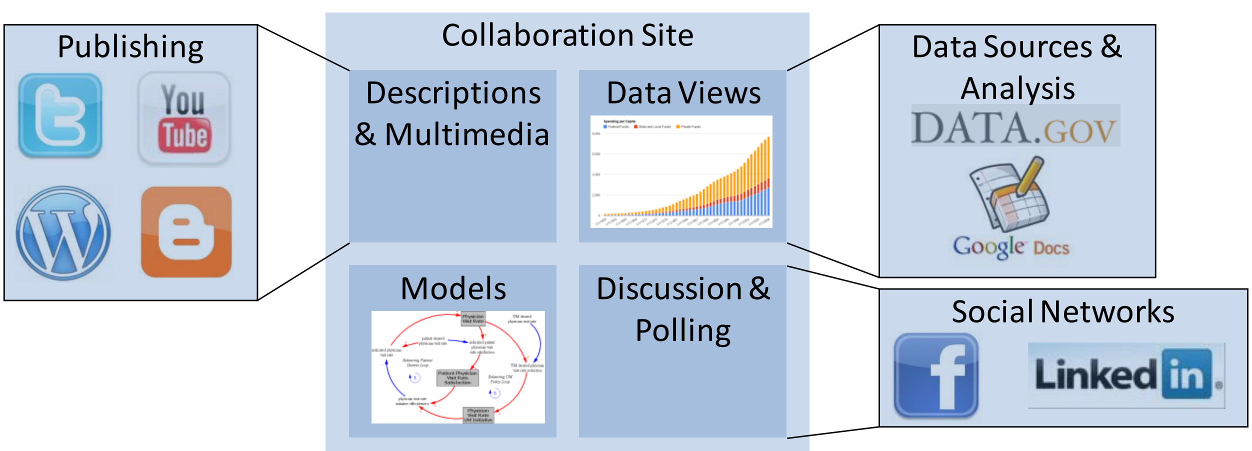

Figure 1: System Overview

Workflow

The design of the workflow for the site is generally an extension of the group model-building process to an online environment that allows users with different expertise to engage at a range of levels of involvement. Model development is not a simple process and will likely always require individuals with significant system dynamics expertise to take part, yet there are also a number of stages where subject matter expertise and stakeholder feedback is essential. Figure 2 provides an overview of some of the key steps in the problem solving process. Some of the activities for each stage include:

- Contributing written and multimedia based problem definitions as well as collecting, analyzing, and charting data for reference modes
- Justifying and collecting data on relevant factors and creating causal loop diagrams
- Developing simulation models and testing and comparing
- Creating gaming simulations and multimedia for stakeholder education

Note that while there is a logical progression from one step to the next, the findings of one stage often require revisiting previous stages. The site will therefore need to help manage this iterative process allowing for version control and branches.

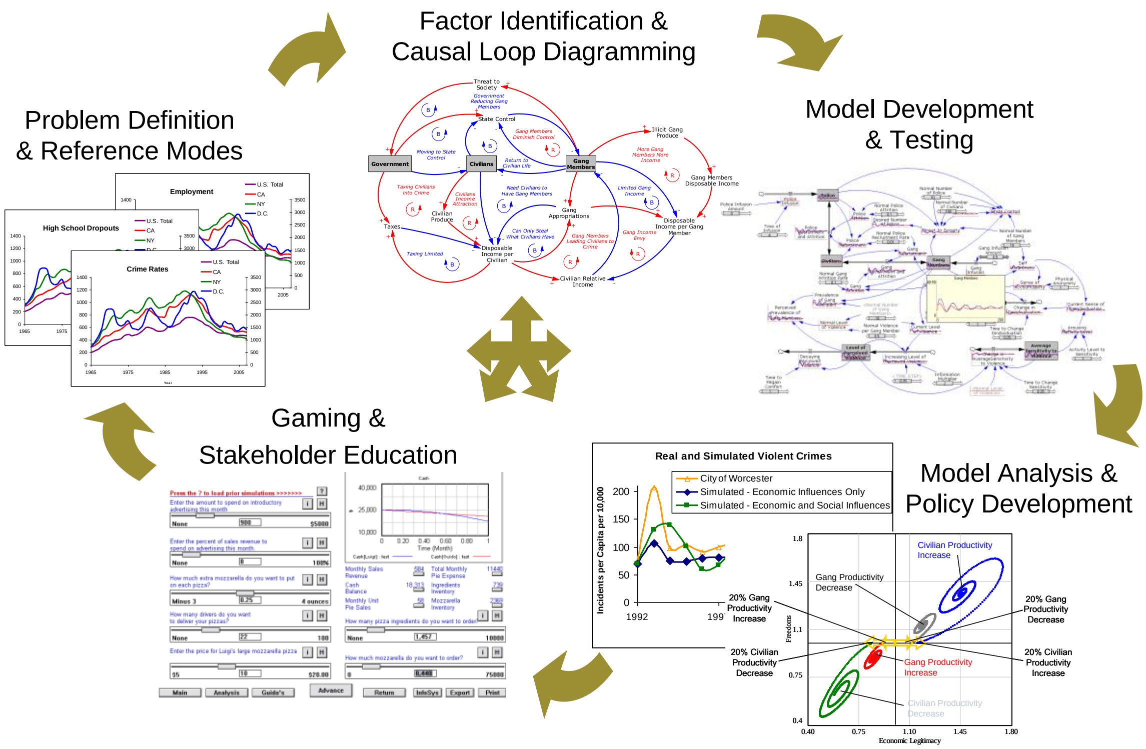

Figure 2: The model development process

Given the sheer size of problems like health care, the design will also need to enable hierarchical relationships across different levels of detail. Through the use of sectors, a more aggregate model of the entire health care system will be related to more detailed descriptions of sub systems. For example, the population model at the highest level might simply be a single stock model while a sub model may use arrays or an agent based representation to study differences across age groups, gender, ethnicity, and income.

Over time, as multiple iterations at different levels of detail are completed, the result will be a comprehensive representation of all the identified issues within a given problem. This representation will have significant advantages over lengthy reports and legislation. Instead of a static snapshot of the problem, this web-based approach will provide a dynamic view of the system that can be easily updated. The use of models will also help with comparing and evaluating the effects of different policies. Instead of relying on clichés and conjecture, such a process will require contributors to provide robust descriptions that can be verified or refuted in quantitative terms.

Contributor management and preventing abuse

The second and third key challenges involve the recruitment and organization of a wide range of users and preventing the manipulation of the system. By leveraging existing "trusted networks," the proposed system will be able to address both of these issues. Large social networks such as Facebook and LinkedIn already provide massive user bases and tools for extending their services. For example, a Facebook application can be created to track an individual's contributions to different project spaces. Another application might allow users to create challenges to provide feedback or retrieve data that they can send to people in their network. Challenges could be as simple as rating explanations and policies to more complex tasks like retrieving data and providing an analysis. For more professional oriented networks like LinkedIn, the descriptions of a users' expertise may be used to help match people to tasks. For example, active users could recruit new experts in a particular field by searching their personal networks. A user's profile might also be used to suggest problems that might be of interest.

In many open forums, a complete "free for all" often results in content that is out of place, redundant, or irrelevant to a particular issue. For this reason the site will need to be organized using roles to control how users can interact with the system and what jobs they will perform. The design seeks to strike a delicate balance between roles that are too strict and limiting, which might discourage participation and suppress innovation with roles that are too loosely defined that create confusion and conflict. Some roles considered include:

- **Manager:** An individual that defines a problem and has the capability to exercise control over other roles and the organization of content. For a given problem there will likely only be a few managers. Depending on the complexity of the problem, however, different sub-problems may be decomposed that are then managed by different individuals.
- **Theorist:** Experts in a particular field that can apply theoretical knowledge and empirical research to help describe and solve a problem. Theorists will take part in the non-technical stages of model development and analysis. It is expected that for a given project fewer than a hundred theorists will be involved.
- **Modeler:** Individuals who have experience in applying mathematical and modeling expertise to develop, test, and report models. Initially modelers will need to participate in the entire process and therefore may also be managers. Over time, as other individuals become more familiar with the process, modelers may only need to be involved with the model development and analysis stages. Actual model development is best done by only a few individuals yet, as with managers, depending on the complexity of the problem, different sub-models may be worked on by other individuals.
- **Analyst:** People in this role will help supply and analyze documents, datasets, and scenarios to initialize models and compare simulations. Analysts will be providing assistance throughout all stages of development, except for perhaps model development and testing. Very often the retrieval and formatting of data and information about a problem can be one of the most time consuming tasks. Several hundred analysts may become involved in a large problem, but would not necessarily need to participate full time.
- **Publisher:** A publisher may simply be a blogger interested in the problem or a person who maintains the more formal model presentations and descriptions that enable other

users to understand the problem and models quickly. Smaller problems may only have tens of publishers, but larger projects could have tens of thousands reporting and presenting content.

- **Contributor:** Individuals in this role provide ratings, polling responses, and brief feedback, as well as nominating the most credible models for further development. The contributor role is expected to be the largest group, with up to millions of people involved with reviewing content, findings, and the suggested policies. User interaction may be as simple as watching a presentation or reading a short description and providing a rating. For more involved projects, contributors may be asked to complete surveys or participate in a structured forum or online town hall meeting. Most users in this category may link to the project or interact through Facebook applications or other widgets that are separate from the core system.

Clearly defining the roles of users in the system should provide a scope of expectation for each user, simplifying the tasks they need to accomplish. In cases where multiple individuals qualify for a given tasks, voting can be used to allow other users to elect the best individual for the job.

Another concept that is frequently used in social media in order to maintain user interest is the idea of an incentive structure. An incentive structure would give users a reason to signup, a reason to keep participating, and a reason to encourage others to join. A simple example of a "ranking" incentive structure would be to have different levels that are accomplishable for each user, and as they participate more often or more usefully, they will increase to a higher status level. Starting off at a "Rookie" level and getting promoted to "Novice" would give users a progression, and could help motivate users to try to get to an "Intermediate" level, etc. Reward badges for achievements such as recruiting and contributions are another option for encouraging engagement.

There are numerous other aspects that should be considered in the design of a system for conducting distributed group problem solving, but what has been discussed in this paper is a sufficient starting point for demonstrating some of the concepts. To illustrate many of the design features discussed, an example web site was created using a study of health care dynamics derived from the work of Jim Thompson (2006 & 2010). The next section will describe the site in greater detail and the example it provides.

Technology demonstration: health care dynamics

Health care has become a problem of increasing concern in the United States and other developed countries as the utilization of services and costs have grown exponentially while the quality of care has seen disproportionately smaller improvements. Despite efforts by health care management organizations and governments, per capita costs have continued to climb at a steady pace. In the summer of 2010 the United States Congress passed the most comprehensive reform since Medicare and Medicaid, but not without controversy. Given the complexity of the health care system, disagreements can be expected, but to truly develop a more comprehensive understanding it is necessary to examine the major components in clear, traceable detail. This makes health care dynamics an ideal initial case for demonstrating some of the distributed group problem solving concepts.

The following sub-sections will discuss simple examples of the work that can be contributed by a variety of users following the workflow described earlier. The examples will only examine health care dynamics at a very high level of aggregation, but the opportunities for further disaggregation and detail are readily apparent. As a proof-of-concept, the demonstration site only explores a portion of the technology needed. At the time of this writing the site has largely remained closed leaving many of the collaboration and sharing features left untested. To see these examples online, please visit: www.socialsolver.com/project-definition/health-care

Problem definition and reference modes

One primary problem with the health care system is that the rate of increase in spending per capita is far greater than inflation and wage increases making it increasingly unaffordable for people in the United States. As a result, Federal funding of health care has increased significantly contributing to long term budget deficit concerns. Below is a chart of the spending per capita by source (U.S. Department of Health & Human Services, 2010) that clearly shows an unsustainable trend with costs nearly doubling over the last decade.

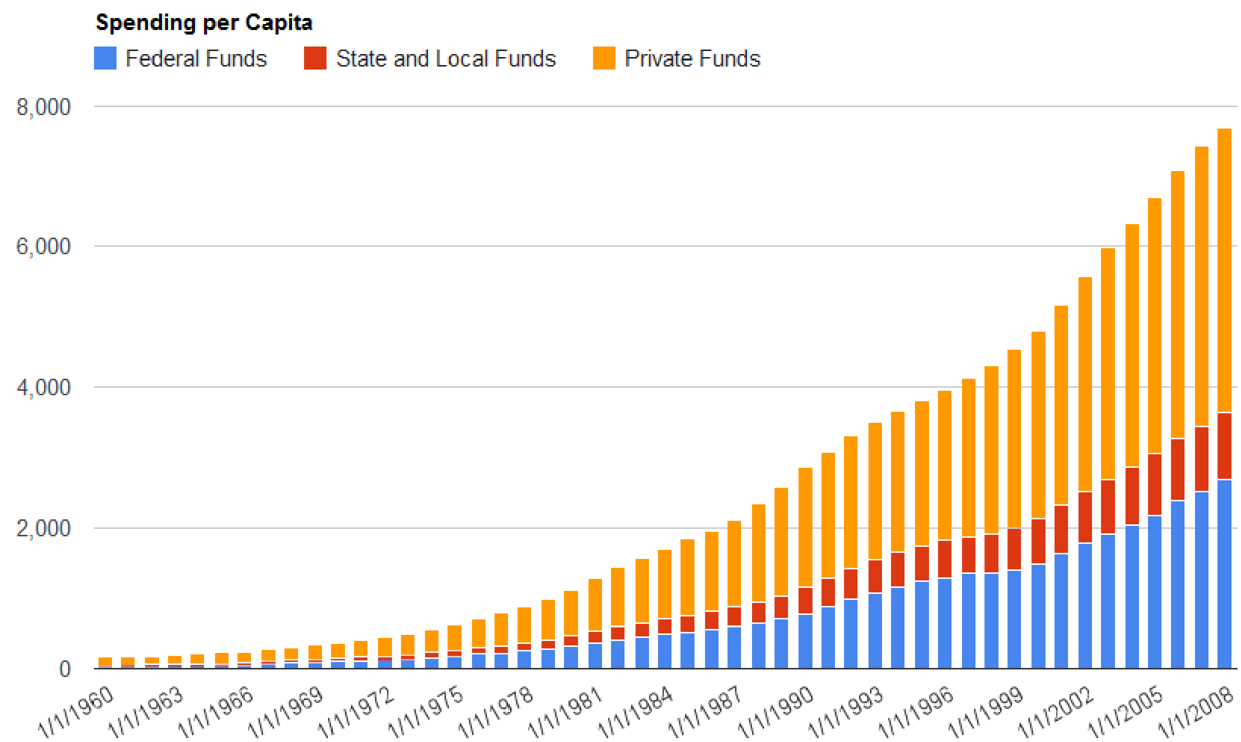

Figure 3: Per capita spending by source

The above chart is actually an embedded version of the one created in a Google Spreadsheet. By using services like Google Spreadsheets, the process of maintaining and updating data is greatly simplified. As new reports are posted, one can simply add it to the spreadsheet where it then automatically updates the associated charts, wherever they may be embedded. Spending is only one of the important reference modes. For instance in this case, another contributor could add the spending data for the years since 2008 to the spreadsheet. Once completed, any page using the spreadsheet would also be updated. There is also likely significant insight in further decomposition of the spending and contributing sources. During this phase of problem description and reference mode sourcing, multiple users would be able collect, clean, and load data and information onto the site.

Factor identification and causal loop diagramming

The next stage of the problem-solving workflow is to identify all the key factors and their causal relationships at a given level of detail. In the developed system, this process would likely lead to many different perspectives on which issues are most important. One benefit of a model-based approach is that these multiple perspectives are all represented in the same format. This allows members to compare different representations and to weigh in on the importance for the current level of detail. Through iteration, factors will likely be added and removed as the group explores the relationships and data. In this example study, the factors only include the sectors of the health care system identified by Thompson (2009 & 2010). It is important to reiterate that this case-study is intended to illustrate the process and tools and not to provide a definitive assessment of the health care system. The sectors loaded on the site include:

- Population growth
- Physician utilization
- Emergency room utilization and capacity
- Outpatient utilization and capacity
- Inpatient utilization and capacity
- Pharmaceuticals
- Medical Technologies
- Physician price setting

For each sector, a set of factors appropriate for the current level of detail are then mapped out to describe the underlying feedback loops. When available, reference mode data on important factors can also be researched and presented. Figure 4 shows a causal loop diagram examining the physician visit rate. Accompanying each diagram is a description of the behavior it represents. In this example, the physician visit rate is moderated by two competing balancing loops. One of the loops is a Utilization Management (UM) policy that seeks to manage the visit rate. In practical terms the UM policy is a very generic structure representing both insurance and government efforts to regulate how often individuals go to a physician. Some example initiatives that have been put into place include co-pays, declaration of a primary care physician, and referrals. If there is a difference with the current visit rate and the UM desired, then an initiative is implemented. This initiative produces incentives that move the visit rate towards a new indicated rate, thus closing the first balancing loop.

The UM policy loop is in direct competition with a second loop that seeks to balance towards the patient desired visit rate. In this case, if there is a gap between the current visit rate and the patient desired visit rate, then the patient satisfaction is affected. The patient satisfaction then modifies the effectiveness of the UM initiative. When patients are 100% satisfied then the policy will also be 100% effective, yet as satisfaction declines, so does the effectiveness of the initiative. The modified initiative then, as before, changes the indicated visit rate, thus closing the second balancing loop.

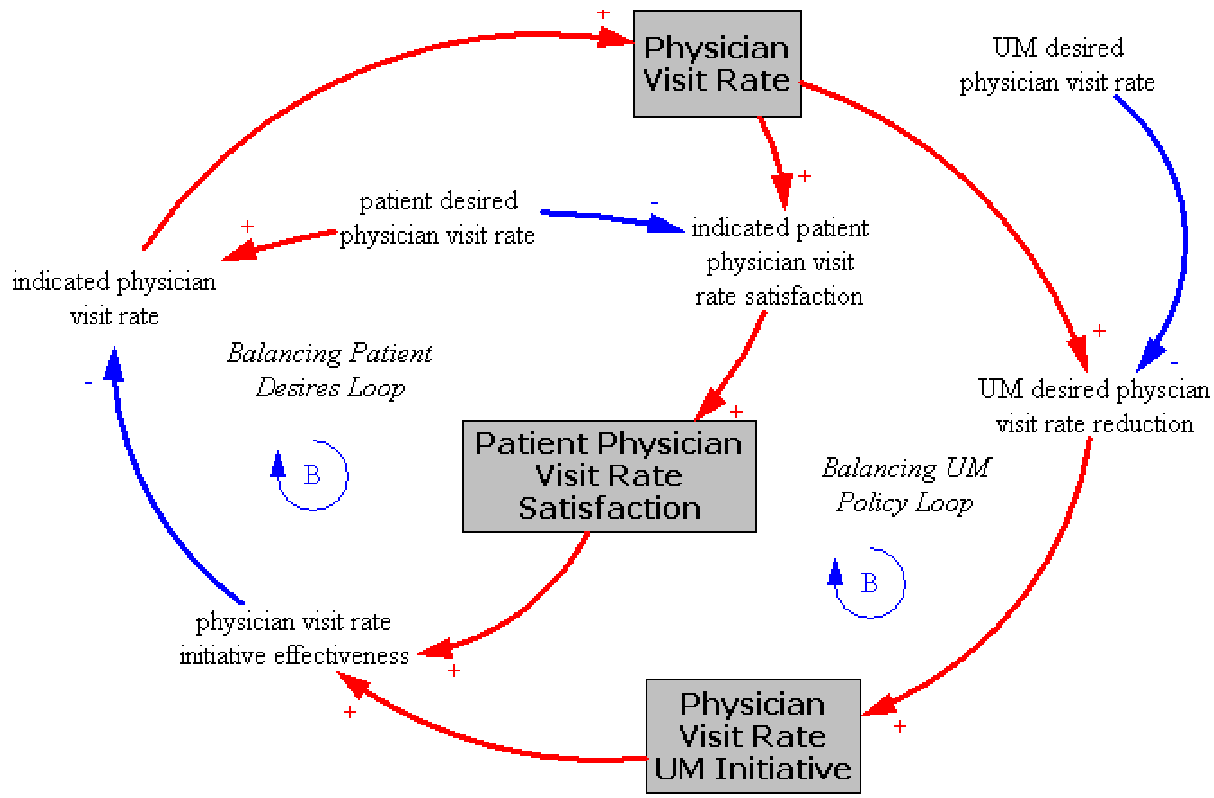

Figure 4: Physician visit rate causal loop diagram

For a reference mode the National Center for Health Statistics (2010) report provides the average physician per capita visit rate shown in Figure 5. For tasks in this stage, both novice and more expert users will be able to contribute. More novice users will be able to help collect and clean relevant data while experts can help explain the possible factors, policies, and changes that may have contributed to the behavior seen in references modes. Working together with a system dynamics practitioner, users will be able to construct causal loop diagrams that are well documented and supported with data.

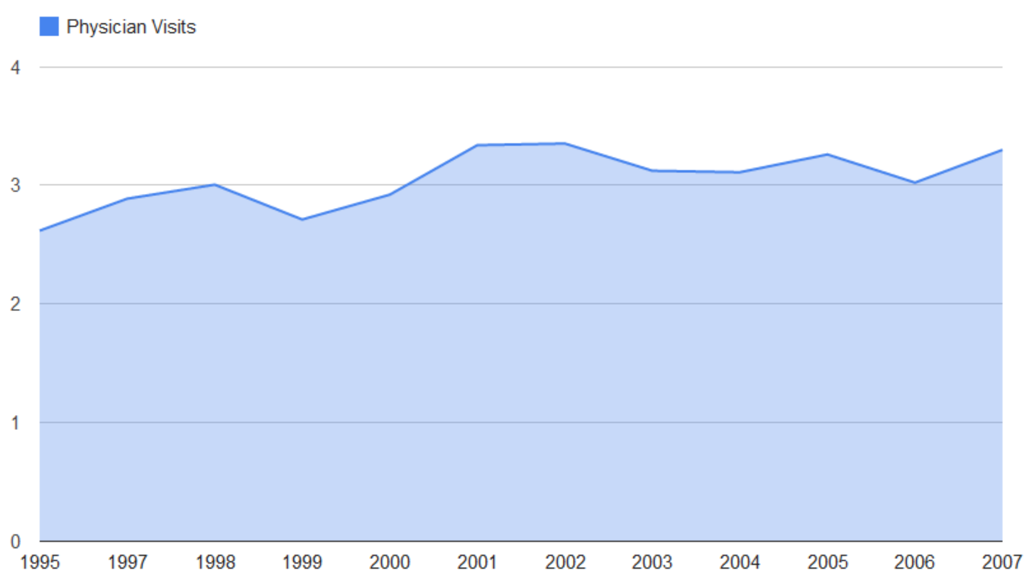

Figure 5: Physician visit rate data

For each of the sectors listed above a similar causal loop diagram was constructed and data was retrieved to provide reference modes for some of the key factors. The information collected in this stage, provides the necessary foundation for the next step of developing and testing simulation models.

Model development and testing

One of the gaps in the services required for the described system is an online modeling tool that is both mature and open. For the purposes of this demonstration, simulation models were created offline in both Vensim and AnyLogic. The AnyLogic models were then embedded in sector pages to allow users to explore how each responded to changes to the equilibrium state. The Vensim models are also attached to each model page. Other options for embedding simulation models include Forio and isee’s NetSim. While embedded models are ideal for many applications, a web based modeling tool would be needed to maintain the openness and accessibility of the proposed system.

Figure 6 shows the model for the physician visit rate shows how the model responds to a drop in the UM desired visit rate in comparison to the historical visit rate. In this model sector only the competing balancing loops shown in the causal diagram are driving the behavior seen. In the simulation, a new UM policy seeking to reduce the annual physician visit rate to 2 from the initial rate of 3. The gap triggers the UM balancing loop, but as shown in the causal diagram it is in competition with the patient desires balancing loop. The end result is a visit rate of approximately 2.8. This behavior is clearly not the same as what is shown in this historical data, but the sector model is not intended to reproduce the reference mode. Instead, this model’s purpose is to help users understand the relationship between the given factors before adding additional inputs. As sectors are later combined into an integrated model the behavior becomes more consistent. Again however, the point of the models, especially in the beginning, is not to find the perfect fit, but rather to develop understanding. With increasing detail and exogenous historical policy inputs a very good fit is certainly possible, but the goal of a project is not to recreate history. The goal is to create a model representation that is descriptive enough that it can help with forming consensus on both the problem description and possible policy solutions.

The contributors at the model development stage will predominately be skilled system dynamics practitioners, but other users will still be able to provide critical feedback on both the clarity of the model and appropriateness of its representation. The models created will serve as the basis for both the policy discussions and the games used for stakeholder education.

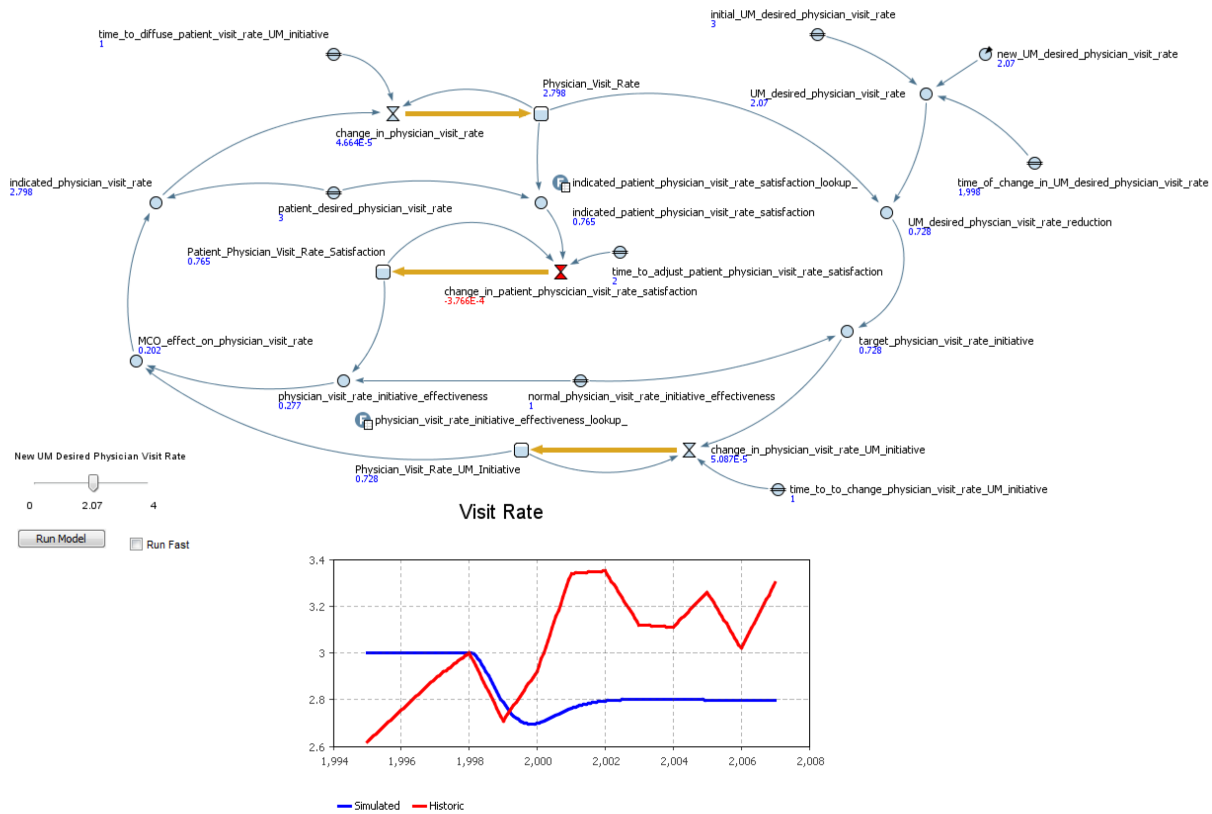

Figure 6: Embedded physician visit rate sector model

Model analysis and policy development

In this demonstration there was insufficient time to conduct any distributed group analysis and policy development, but the insights of Thompson (2006 & 2010) are presented on the site. Some of the findings include the following:

- Much of the health care system is structured to increase costs with pharmaceuticals and medical technology serving as some of the stronger reinforcing loops
- An inpatient's average length of stay has dropped significantly due to pharmaceuticals and medical technology
- These same advances have also encouraged more outpatient visits
- The physician population is close to equilibrium, which does not track with the technology fueled utilization growth

This stage is arguable one of the most important since it where users will be discovering potential solutions. It is also the part most likely to be targeted for manipulation. In this example, pharmaceutical and medical technology companies may be adversely affected by policies that attempt to limit their impact of increased costs. For instance in 2007, pharmaceutical companies spent a reported $168 million in lobbying according to The Center for Public Integrity (Ismail, 2008). While models and public data will help to clarify the structure of the problems with health care, such powerful influence will need to be carefully guarded against. One of the controls discussed earlier is the use of trusted networks. While not impervious to exploitation, networks like Linked-In and Facebook have invested significant resources to ensuring that their users are not spammed or recipients of unwanted solicitations. Other measures such as inter rate reliability could be used to identify questionable individuals and groups.

As with the previous stage, more complex model analysis will require more seasoned modelers and policy experts. Yet a broader range of users can provide important feedback by interacting with the models, documenting their behaviors, and discussing the proposed policy options.

Gaming and stakeholder education

The final step of the process is to begin producing materials and gaming simulations that allow a wide ranging audience to begin learning from the findings of a project. Building off of the models developed earlier and the policies proposed, the goal of this stage will be to make the key concepts understandable to as large an audience as possible. Possible materials include videos and presentations that summarize the different model sectors, games that challenge individuals to manage the health care system, and polls that help to assess understanding of the problem and areas for future more detailed studies. Some example of simulation-based games for grand challenges include C-Learn (2011), which helps users to understand the long-term impact of CO2 levels, and HealthBound (2011), which explores policy tradeoffs for addressing health care quality and costs.

As discussed in the introduction, one of factor that defines a grand challenge like health care is its deep entanglement with the behaviors of millions of people. It is likely that in order for there to be measureable success, individuals will have to agree to make perhaps difficult changes that cannot simply be mandated by law. For example, the government mandate in the current health care legislation calling for every individual to purchase health care has already been challenged in courts. Policies that are outside of the traditional top down control therefore require a bottom up social movement that is built on education and outreach. Social networks like Facebook and Twitter have already demonstrated their potential in this regard, yet whether or not they can be applied in such a constructive process is still to be determined.

Conclusions

The goal of this paper was to present a new framework and workflow for taking on problems that defy more conventional problem solving approaches. These “grand challenges” are made particularly difficult because of their deep entanglement with social issues. It is argued that system dynamics can be applied to help develop a new process by combining its mature problem solving approach with powerful internet technologies and social networking tools. The resulting system will be one that enables users with wide-ranging levels of expertise to engage in the process of understanding and developing solutions for grand challenges.

Through a demonstration study of health care dynamics, this paper sought to illustrate how existing technologies can be integrated to provide much of the necessary functionality. Some of the technology areas requiring additional development include online model building tools, methods for effectively organizing hierarchical content, and a system for commenting and rating content. Further design consideration should also be made of methods for preventing the abuse and manipulation of the system. In addition, mechanisms for recruiting and incentivizing a massive number of users are needed.

The next steps planned for advancing the concepts presented in this paper are to open the site to a larger group of users with different levels of expertise and to initiate additional example grand challenge studies.

References

CDC HealthBound Game (2011). Retrieved April 9, 2011 from http://www.cdc.gov/healthbound/

C-Learn (2011). Retrieved April 9, 2011 from http://www.cdc.gov/healthbound/

Ismail M.(2008). “A Record Year for the Pharmaceutical Lobby in '07,” The Center for Public Integrity. Retrieved April 9, 2011 from http://projects.publicintegrity.org/rx/report.aspx?aid=985

National Center for Health Statistics (2010). “Health, United States, 2009: With Special Feature on Medical Technology.” Hyattsville, MD.

Thompson J . (2006). “Making Sense of US Health Care System Dynamics,” Proceedings of the 2006 International System Dynamics Conference, Nijmegen The Netherlands.

Thompson J. (2010). “Making Sense of U.S. Health Care Macrodynamics: Medical Technology, Freedom of Choice and the Pursuit of Immortality”, Worcester Polytechnic Institute Social Science and Policy Studies Lecture Series, Worcester MA

U.S. Department of Health & Human Services, 2010. “National Health Expenditure Data”. Retrieved October 27, 2010 from http://www.cms.gov/NationalHealthExpendData/02_NationalHealthAccountsHistorical.asp