COMPLEX, SOCIOTECHNICAL SYSTEMS (CSS): SOME FUNDAMENTAL CONCEPTS

Joseph Sussman   February 5, 2011

Teaching note: ESD.00, Introduction to Engineering Systems

Introduction and Proviso

Our interests in ESD.00 are largely focused on complex, sociotechnical systems (CSS)—understanding and improving them and developing methods for representing and designing them. By CSS, we mean (for our purposes here), large-impact, difficult-to-understand systems that have important technology issues *and* important societal implications. We will spare you, for now, the myriad definitions of "complexity".

Throughout this note we will refer to CSS, but we recognize that much of what we say about CSS is true of other systems too—technical systems, social systems and others as well.

This note attempts to capture many concepts that we experientially have found to be of value when one tries to characterize/analyze/design/think about a CSS. No claim is made for completeness; there are doubtless other valuable concepts. Nor, except in a few situations, do we claim originality; many of these concepts will be familiar to you and have been documented in texts, professional papers and so forth. We owe a huge debt to seminal systems authors including (but not limited to) Forrester, Simon, Senge, Lorenz, Weiner and others. Also we have drawn on the work of ESD colleagues including (but not limited to) Sterman, Moses, Magee, de Weck and others.

We now introduce a number of CSS concepts. We will use the Boston/Cambridge transportation system as a CSS example to illustrate some of the concepts.

CSS Concepts

When we study CSS, we need to take a broad view of the problem space, recognizing there are interactions of **"our"** CSS with other systems external to it that often need to be explicitly considered. This concept can be thought of as "horizontal" in that while considering our CSS, we should consider other parallel systems to which it is connected.

Our CSS of interest is always part of a larger system composed of our CSS and these parallel systems; our CSS affects the behavior of the parallel systems and hence the larger system; in turn our CSS's behavior is affected by the parallel systems and the larger system.

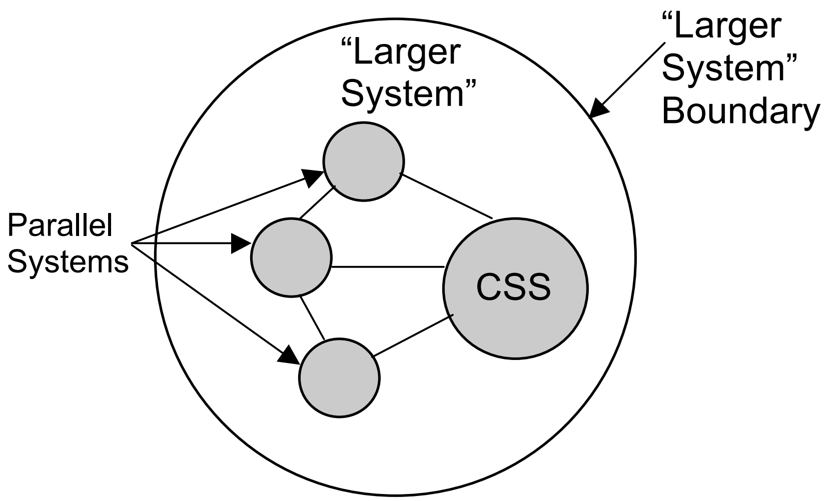

Some like to think as the set of parallel subsystems as the "environment" of the CSS and we will do that in this document.

Our CSS, the Boston/Cambridge transportation system, has various parallel systems. These might include:

- *the air quality in Boston/Cambridge,*
- *MIT itself (whose behavior will affect the Boston/Cambridge transportation system: when MIT is on spring break, the loads on the transportation system are certainly lower),*
- *the commuter rail system which provides service from the outlying suburbs to Boston/Cambridge (commuter trains provide service to both North and South Stations)*
- *the economic system in Boston/Cambridge, clearly affected by the effective delivery of goods, the ability of people to get to work or school and the vibrancy of the economy clearly relates to the loads on the transportation system*
- *the real estate market in Boston/Cambridge, will relate to accessibility of a residence to job opportunities.*

It follows that our CSS is composed, in part, of smaller systems—often referred to as "subsystems".

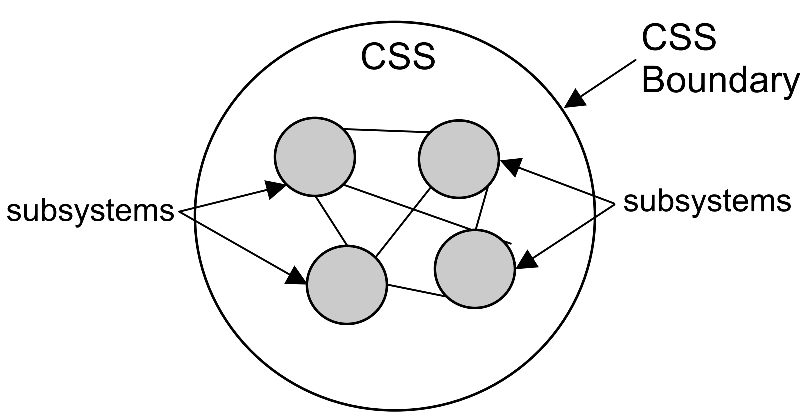

Of course, these two diagrams represent the same concept at different levels of hierarchy.

Subsystems of our CSS include:

- *The highway network*
- *The MBTA system, which would in turn include vehicles buses and trains, and the infrastructure needed for the MBTA to operate which includes the rail network, the highway network on which the MBTA's buses operate, fare collection systems, operating staffs*
- *Travelers who will make a choice of what mode to use: driving, taking public transportation of one sort or another, traveling intermodally (bike to the station and take the train from there)*

The micro-macro question

A key systems issue is the relationship between the macro-behavior of our CSS, and the micro-behavior of subsystems. It is important to understand how that micro-behavior affects macro-behavior of our CSS, and at what level-of-detail we have to understand the micro-behavior in order to properly understand the macro-behavior. Further, we need to understand how the subsystems affect each other and how those interactions impact on the macro-behavior of the CSS. So perhaps there are rules of thumb about how deeply we need understand the micro and its relationship to the macro—but in practice, we go up and down the chain, down into the micro (the subsystems), back up to the macro, and so on as one iterates until we find out what we "want to know" with some level of confidence, which of course will vary from person to person.

In our CSS, do we need to know how an internal combustion engine works? Probably not. Do we have to know if a street has bike lanes? Probably. But it depends upon the questions we want to answer about our CSS.

For example, we may be interested in the electric power grid, at a macro-scale. So, this is our CSS. The behavior of this CSS is affected by user behavior at a micro-scale, such as the household use of devices like an "energy box" to optimize household energy costs. At what level-of-detail do we need to understand how the energy box affects household use in order to have a good understanding of the behavior of the CSS of interest, the electric power grid? Similarly, thinking "horizontally", the interaction between our CSS and the regional environmental CSS may be of interest as well.

We need to recognize that the optimizing of each of the subsystems of a CSS may lead to overall CSS performance that is far from optimal. *"Everyone"* understands this, but in real life, it is a huge temptation to do exactly that and sub-optimize your CSS! Of course, not trying to optimize even the subsystems—also a temptation when faced with all this complexity-- likely will lead to even poorer CSS performance.

Feedback

So put another way, we need to understand the implications and impacts of *feedback* among subsystems internal to the CSS, and between the system environment—made up of the parallel systems-- and the CSS itself. Also we must recognize the *time delays* inherent in feedback and their effect on CSS behavior.

For example, there is feedback between the housing location choice of individuals and the quality and connectivity of the transportation system in Boston/Cambridge. Shape the CSS is a particular way and people will settle in particular locations. But you may shape the transportation system to accommodate new settlement patterns.

Emergent behavior

In our consideration of CSS, we need to train ourselves to look for emergent behavior—new behavior that arises in CSSs as a result of interactions at the subsystem level -- and unanticipated consequences characteristic of these CSSs. These behaviors and consequences can be "good" or "bad". It seems though that "bad" occurs more than "good"!

It is always hard to come up with good examples of emergent properties, because if you can identify them in advance, they may be all that interesting. An example here could be the linkage of public health and the transportation system. We now recognize that species diversity is linked to air and water quality that is in turn linked to the transportation system—how many people choose to drive. But it was not so long ago that this linkage wasn't recognized, or least not worried about.

Non-linearity

We must recognize the implications and impacts on CSS behavior of *non-linearity*, in subsystem behavior internal to the CSS, and between the CSS environment—made up of the parallel systems--and the CSS itself and in feedback that may be non-linear in nature.

For example, the travel time on a road segment is a highly nonlinear function as the volume on the road approaches the roadway's capacity

Uncertainty

We must consider the implications and impacts on CSS behavior of *uncertainty* in subsystem behavior internal to the CSS, and between the CSS environment—made up of the parallel systems-- and the CSS itself and in feedback that may be stochastic (i.e. uncertain) in nature. We are often concerned about uncertainty of CSSs, particularly when we are attempting to make predictions of future events (e.g. climate change, transportation demand).

There are any number of uncertainties in our CSS. For example, the weather clearly affects the performance of our CSS. Thinking longer term, the state of the economy, which is surely uncertain, will clear affect the loads our CSS will be expected to deal with.

Cause and Effect in CSS

Cause and effect may differ in time: occur at different times
Cause and effect may differ in space: be geographically distant from each other
Cause and effect may differ in kind—e.g. changes in the transportation system causing health effects (we know this now, but didn't always know it—what connections between cause and effect are we unaware of now that will bite us someday?)

It may be difficult to understand the directionality of cause and effect—e.g. in developing countries, is the rise in income causing the fall in infant mortality or does the fall in infant mortality cause a rise in income? Or e.g. does transportation investment drive land use or does land use change lead to transportation investment or perhaps are both true?

Effects may arise from multiple causes; further, causes may result in many effects. In considering cause and effect, it is important not to confuse correlation and causality.

Small changes in initial conditions or in a parameter can lead to very large differences in CSS behavior over time.

Long-term vs. Short-term behavior

There are usually considerable differences between the long-term and short-term behavior of CSSs. Positive short-term behavior, in response to an intervention we implement, may give us false confidence; this positive behavior may be followed by negative behavior in the long-term.

Time scales

Components of CSSs may operate at very different time scales. For example, in transportation systems, operating policies may be changed in the short-term; vehicle purchases may occur in the medium-term; and building infrastructure will take place in the long-term.

Humility

The points above create considerable difficulties in predicting system behavior; when we make design decisions that affect that behavior, it requires humility on our part. Often we are happy simply to be able to predict "the sign" of system change in response to actions we might take. It is clear that for successful design, we would like to predict the behavior fairly accurately but at present this may not be possible. While this does not suggest that we do nothing, we must try to understand the characteristics and behavior of a CSS, as well as we can.

Evaluative Complexity

Multiple stakeholders of complex, sociotechnical systems have their disparate interests and hence their different ways of evaluating system performance. We call this *evaluative complexity*. Resolving evaluative complexity often creates ethical dilemmas as, in effect, choices among stakeholders are made by the system designers and operators.

Raising bus fares to build capital to buy more buses may be favored by the MBTA, but travelers might think otherwise. What might be MIT's opinion of higher bus fares?

Rationality

People and organizations are *not* always "rational" in the way they make decisions and hence the initiatives we implement as an intervention may not produce the behavior we expect if we assume "rationality".

Further, there are limits to human capability to understand a CSS. This is an illustration of what Herbert A. Simon calls "bounded rationality".

People and organizations respond to incentives (e.g. through pricing) or at least how they perceive their incentives. This can be a means of altering their behavior and hence overall system behavior. But as noted, we may not see their behavior as "rational".

One might think someone irrational if they live on a bus route in Cambridge, but choose to drive to work everyday. You wonder why that individual sits in rush-hour traffic everyday. But the decision calculus of the person may weigh other factors (convenience in dropping off and picking up your child at day care) more heavily. In the US, the number of people in car pools is half that is was in 1980. Think about why that may be.

Network structure and beyond

We often find CSSs are structured as networks—physical, logical and organizational—with flows on the links of these networks of matter, energy and information, traveling between nodes.

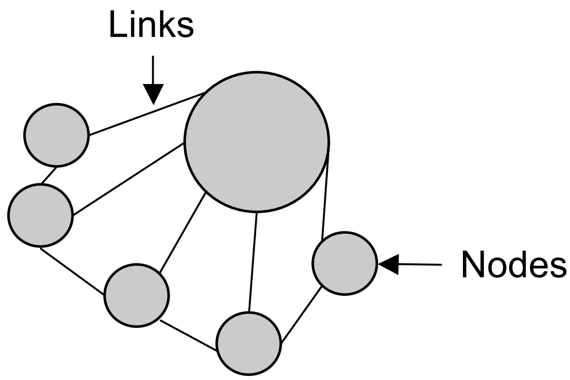

Performance of network-based systems is affected by the behavior of the links and nodes. Nodes may be more critical in some cases and links in others, depending of the type of network. For example, in rail freight networks, experience has taught us that node (terminal) performance dominates overall system performance.

One could argue that transportation systems are the original network. Clearly the way the links and nodes are interconnected determine the performance of the Boston/Cambridge transportation system. The amount of parking MIT (a node) provides relates to the number of people who choose to drive and hence the flows on the highway links.

CSS Design

Sustainability

Sustainability--- defined by the three Es, **E**conomic development, concern with the **E**nvironment and social **E**quity--- is the overarching design principle for complex, sociotechnical systems (simply a Sussman assertion, I freely admit).

Goals, Objectives and Performance Measures

The design process requires the development of system goals and objectives as well as performance measures.

It can be surprisingly hard to choose goals, and objectives for a CSS in such a way that one can make decisions about what to do. And performance measures will be different for the operator of the bus system (who is concerned with costs as well as the quality of the service provided) and bus travelers who are concerned with getting somewhere on time at a reasonable fare.

Strategic Alternatives and Robust Bundles

The creative aspect of system *design* is the generation of individual *"strategic alternatives"* dealing with decisions about technology, operating practices, investment, organizational change and changes in relationships among organizations, and so forth.

Strategic alternatives for our CSS are numerous. Buy more buses, add more bike lanes, change signal timing are but a few. And they can be organizational too. For example, one might (try to) integrate MBTA operations with the MIT on-campus shuttle.

And because there is rarely a "sliver bullet", one strategic alternative that solves all our problems, we usually need to implement bundles of strategic alternatives that are "robust" in that the bundles will work well for a wide range of conditions.

CSS evolution

Further, as we design, we must recognize that complex, sociotechnical systems may change of their own accord during the design process; they evolve and are never are static.

For example, we are designing our CSS, the Boston/Cambridge transportation system and meanwhile, MIT is developing its campus in new ways or is expanding its number of undergraduate and graduate students.

Considering External Response

A common error in CSS design is that we do not consider the *external response adequately (competition, customers, regulators, suppliers, etc.)*. We often do not understand the change in expectations of customers or others external to the system, based on system performance and level-of-service provided. For example, we often have a static assumption about what our competition will do when we introduce new concepts; if the US builds high speed rail in regional corridors, we can't expect the airlines to simply watch it happen without reacting.

MIT has to be concerned about the reaction of the Cambridge community (and beyond) when it builds more parking structures. Might that lead to new regulations? The MBTA has to recognize that raising fares will induce different traveler behavior.

*"*Satisficing*"*

In CSS design, optimization is often beyond the pale. We are happy enough to find a few feasible "solutions", never mind optimal ones. We search for alternatives that are "good enough" -- "satisficing" -- especially when there is limited analysis capability or limited information or limited time.

Conclusion

We hope you find this list of "systems concepts" of some value; perhaps it can help guide you in the way you think about systems. Later this semester, when you have become more expert in thinking about systems, we will provide some additional aspects of complex, sociotechnical systems.