On Theory-Driven Design of Collaboration Technology and Process

Robert O. Briggs

Delft University of Technology, University of Arizona
bbriggs@groupsystems.com

Abstract. The design and deployment of collaboration technology has, until lately been more of an art than a science, but it has produced some solid successes. Commercial groupware products now support millions of collaborations per year. Under certain circumstances teams that use Group Support Systems perform far better than groups that do not. However, as impressive as the achievements are in this field, we can do better. A rigorous theoretical approach to the design of collaboration technology and process can lead us to non-intuitive design choices that produce successes beyond those possible with a seat-of-the-pants approach. This paper explains the simple structure of a rigorous scientific theory and offers examples of theory-driven design choices that produced substantial benefits. It then differentiates rigorous theory from several classes of theory that have intuitive appeal, but cannot inform design choices. It then argues that the logic of the theory-driven design approach suggests that the most useful focus for collaboration technology researchers would be the technology-supported work process, rather than just the technology.

1 Collaboration Technology Design as an Art

Designing collaboration technology has, until lately, been an art, founded on common sense and intelligence, guided by heuristics derived from inspiration tempered by hard experience. This approach has given rise to some solid long-term successes – consider, for example Lotus Notes, NetMeeting, and Webex, each of which now supports millions of collaborations per year. A robust body of literature shows that, under certain circumstances, people who use Group Support Systems (GSS) can be substantially more productive than people who do not (see Fjermestad and Hiltz, 1999, 2001 for compendia of GSS lab and field research). In 1999, we surveyed 127 organizations that used GSS. They reported an average cost saving of $1.7 million per year on an average technology investment of $75,000 USD. A year-long field study of more than 60 groups of GSS users at Boeing (Post, 1992) showed an ROI of 689% on an investment of approximately $100,000. It is rare in any field to find returns at that level.

Such results are nothing short of spectacular, yet as good as they are, we can do better; much better. Brilliant, intuitive minds and seat-of-the-pants reasoning can only carry us so far.

As many successes as there have been in this field, there have been many more failures. Remember The Coordinator? In the early 1990's, this system was heavily funded and highly touted. It was one of the first attempts at an integrated virtual

G.-J. de Vreede et al. (Eds.): CRIWG 2004, LNCS 3198, pp. 1–15, 2004.

workspace, including team calendaring, document repository, and a handful of other seemingly useful technologies. Yet users hated the system. Some disparaged it as Nazi-ware because of the draconian patterns of collaboration it enforced (e.g. every inbound message shall receive a reply before other work could be continued in the system).

The Boeing GSS case was so successful that it was written up in the February, 1992 issue of Fortune Magazine. And yet, the same week the article appeared, Boeing disbanded their GSS facility and reassigned all its personnel to other projects. This pattern of significant success followed by sudden cessation has been repeated by many organizations that adopt GSS (Agres, Vreede and Briggs, 2004, Briggs, Vreede and Nunamaker, 2003).

In light of these examples, consider these questions:

- How can we account for the dramatic success of some collaboration technologies?
- More importantly, how can we repeat those successes elsewhere?
- As successful as some collaboration technologies have been, are they as successful as they could be?
- How would we know?
- What else could we try that has never been considered that would be all but guaranteed to work?
- How can we account for stunning failures of other collaboration technologies?
- More importantly, how can we avoid them elsewhere?
- Do we have to wait for inspiration to strike some genius in our field before we attain our next leap forward?

Good theory can address all these questions.

2 There Is Nothing So Useful as a Good Theory

There is nothing as useful as a good theory. This assertion may draw snorts of derision from skeptics who may regard theory as an excuse for not doing anything useful. Yet, a good theory can put people on the moon and return them safely to earth on the first try. What one theory can do for space travel, others can do and have done for collaboration technology. Rigorous theory can lead to designs for collaboration technology process that far surpass those produced by a good mind and a gut feel. This paper explains what is meant by *good theory*, and present several cases to illustrate how a good theory can drive the design and deployment of collaboration technology in non-intuitive ways to yield unexpected success. It also discusses several classes of theory that seem tantalizingly useful at first, but which cannot lead to useful results. It concludes by discussing implications for ongoing collaboration technology research.

2.1 Good Theory: Always a Causal Model

A good scientific theory is a *model* of *cause-and-effect* to explain some *phenomenon of interest*.

Every technology presumes a cause-and-effect. Every technology is built to improve some outcome. By definition, this presumes that mechanisms exist to cause changes in the outcome of interest, and that technology can be used to invoke those mechanisms. A theory is a model of those causal mechanisms. The theory gives us a basis for understanding how we might use technology to attain the outcomes we want. If we have not taken the time to rigorously articulate the mechanisms that cause the outcome of interest, our technologies and our processes may miss the mark.

2.2 Phenomenon-of-Interest: Always the Effect; Never the Cause

In a good theory, the phenomenon-of-interest is *always* the effect; the outcome. The phenomenon of interest is the unchanging focus of a theory. The first step to successful theory-driven design is to explicitly identify, and then define the phenomenon of interest. It is not always obvious at first what outcomes a technology is meant to improve. For collaboration researchers, the possibilities are many – productivity, creativity, satisfaction, and so on. Suppose a technology designer wanted to improve satisfaction with group processes and with group products among stakeholders who were negotiating requirements for a new software development project. The phenomenon of interest would be satisfaction. It would not be group process, not group product, not requirements negotiation, not software development, and not project management. The research question would be, “What causes people to feel satisfied?”

However, labeling the phenomenon of interest is not sufficient. It must be explicitly defined. For example, the word, satisfaction, has many connotations in the English language. In one sense, satisfaction could be a *judgment* that goals have been attained or constraints have been met. In another sense, satisfaction could be an *emotional response* pertaining to goal attainment. These different satisfactions spring from different causes, and so have different theoretical explanations. Theory-driving design must therefore begin by not only identifying and labeling, but also by explicitly defining the phenomenon-of-interest.

Having identified and defined a phenomenon of interest, the next step is to challenge one’s judgment, asking whether this outcome is truly the most useful or important target for improvement. If this outcome were improved, whose work might be more effective? Whose life might improve? Is there other outcome that could be improved instead to yield better results? When one can present an unbreakable case that a certain outcome is truly worthy of effort, it is then time to seek good theory.

2.3 Good Theory: Constructs Connected by Propositions

A good theory is a model of cause-and-effect that can account for variations in the outcome of interest. The logic of these models can be represented as collection of statements with a particular structure. These statements have only two components: axioms and propositions. For example:

If we assume that:

(Axiom 1) all individual actions are purposeful toward attaining goals,

Then it must be that:

(Proposition 1) the effort an individual expends toward attaining a group goal will be a function of goal congruence (the degree to which the group goal is compatible with the individual’s private goals.)

An axiom is nothing more than an assumption about some mechanism that could affect the phenomenon of interest. An axiom doesn't assert Truth (with a capital T). It simply prompts one to ask, "If this assumption *were* true, would that be sufficient to explain how to get the outcome we want?"

A proposition is a functional statement of cause and effect. A proposition posits a causal relationship between two constructs. Constructs are different than variables. A variable is measurable. Productivity in the context of an automobile factory can be measured by the variable, "number-of-cars." Productivity in a brainstorming group can be measured in terms of the variable "number-of-ideas." But productivity itself is not a variable, it is an idea.

Propositions always posit that one construct causes another. A construct in a proposition is either a cause or an effect. In Proposition 1 above, the constructs are Effort and Goal Congruence.

There are a number of different ways a given proposition could be stated. The examples below express the same construct with different words:

- Individual effort toward a group goal is a function of goal congruence
- Goal congruence causes individual effort toward a group goal
- Individual Effort is determined by goal congruence
- The more goal congruence, the more effort.

All these phrasings can be summarized by a mathematical function like this:

$$E = f(G) \tag{1}$$

Where:
E = Individual Effort
G = Goal Congruence

Notice that this very simple expression is not specific about the nature of the function; it does not specify whether it linear, curvilinear, or discontiguous; it does not specify whether the function is bounded or infinite; whether it has constant or variable parameters. Such a young, incomplete theory would no doubt acquire more nuance and complexity as research progressed. Nonetheless, simple though it is, it is still useful to the collaboration technology designer. It suggests that E is a positive function of G, which means if you can figure a way to use technology to increase G, you should get more E as a result.

Theoretical propositions can also be illustrated as simple box-and-arrow diagrams. Figure 1 illustrates three propositions in a simple theory of group productivity. Each box-arrow-box combination constitutes a proposition. The direction of the arrow indicates the direction of causation.

In Figure 1, notice that Proposition 2, the Distraction proposition, posits an inverse relationship rather than a positive relationship. It could be interpreted, "The more distraction a group experiences, the less effort it will make toward a goal." Thus, if you could find a way to use technology to reduce distraction, effort toward the goal should increase, which should in turn increase productivity.

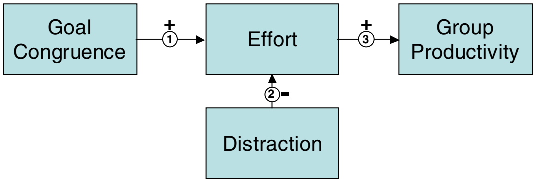

Fig. 1. A box and arrow diagram of several theoretical propositions to explain group productivity. The model posits Productivity as a positive function of goal-directed Effort. It posits Effort as a positive function of goal congruence, and as a negative function of distraction.

Box-and-arrow diagrams provide a way to run the Absurdity Test, a quick, useful way to smoke-test the basic logic of a proposition. Box-arrow-box combinations for improperly framed propositions will yield absurd statements when they are interpreted in the following form:

"The more of <Box X> we have, the more of <Box Y> will result."

For example, consider the two propositions in Figure 2.

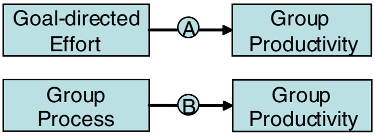

Fig. 2. Two Theoretical Propositions. Proposition A passes the Absurdity Test. Proposition B fails the test.

Proposition A yields a sensible statement when put in the form, "The more goal-directed effort group members make, the more productive the group will be." On the other hand, Proposition B, which has some intuitive appeal, (Group productivity must surely be a function of group process), nonetheless yields an absurd statement when put into the form, "The more process a group has, the more productive the group will be." The Absurdity Test quickly reveals a logical flaw in the proposition, signaling the need for additional attention.

Notice also that the axioms do not appear in the box and arrow diagram. Nonetheless, every proposition is based on one or more assumptions, whether or not they have been articulated. Until the axioms have been teased out and explicitly articulated, their validity cannot be judged, and so the model does not yet fully explain the phenomenon of interest, and the theory is not yet complete.

To summarize, then, a causal theory is a collection of statements that propose mechanisms that could cause a phenomenon of interest. These statements are composed of axioms (assumptions) and propositions (functional statements of cause and effect). They combine to form the logic of a causal theory, making arguments that take the form, "If we assume X, then it must be that Y is a function of Z. The propositions of a causal theory can be illustrated with a box-and-arrow diagram. Each box-arrow-box combination can be interpreted as some variation on the theme, "The more Z you have, the more Y will result."

3 Good Theories – Better Technologies

This section presents three examples that illustrate how a good theory can drive non-intuitive design choices that improve group outcomes.

3.1 Focus Theory and the Brainstorming Feedback Graph

The first example began with work started more than 50 years ago. In 1953, Osborn proposed a new group ideation technique that he called brainstorming. He conjectured that ideation could be improved if people followed a four-rule protocol:

- Do not criticize other people's ideas.
- Be open to wild or unusual ideas.
- Generate as many ideas as you can.
- Build and expand on other people's ideas.

Osborn's reasoning seemed sound, yet twenty subsequent studies were not able to demonstrate that people using the brainstorming protocol produced more ideas than nominal groups (Diehl & Stroebe, 1987; Valacich, Dennis, & Connolly, 1994).

Diehl and Stroebe (1987, 1991) unraveled the mystery by demonstrating that brainstorming groups suffer from production blocking, evaluation apprehension, and free riding. A number of subsequent ideation studies demonstrated that production blocking and evaluation apprehension could be overcome by using Group Support Systems (GSS) that allowed participants to contribute their ideas simultaneously and anonymously over a computer network. People using GSS could all contribute to a brainstorm simultaneously, which eliminated production blocking. People using GSS could also contribute to a brainstorm anonymously, which eliminated evaluation apprehension. The results were dramatic; under certain circumstances, people using GSS produced thirty to fifty percent more unique ideas than did people using nominal group technique (e.g., Dennis and Valacich, 1993; Gallupe, Bastianutti, & Cooper, 1991; Gallupe, Dennis, Cooper, Valacich, & Bastianutti, 1992; Valacich, Dennis, & Connolly, 1994). As remarkable as those results appeared to be, the question remained whether they were as good as they could be.

More than one hundred years of social loafing research showed unequivocally that, regardless of task, people who were working anonymously tended to make less effort than people whose contributions were individually identifiable (e.g. Harkins and Jackson, 1985, Kerr and Braun, 1981, 1983). GSS users were working anonymously, which meant social loafing had to be occurring. The question was whether anything could be done about it.

The electronic brainstorming system we used in our research included a feedback graph (Figure 3 (left)). It plotted the cumulative number of contributions the team made to the brainstorm over time. However, research had not shown that use of the graph had any impact on brainstorming productivity. The Focus Theory of group productivity (Briggs, 1994) suggests that effort toward the group goal is a function of goal congruence (the degree to which the private goals of individual members are compatible with the public goal of the group). Social Comparison Theory (Goethels and Darley, 1987) suggested a goal congruence hook that could be invoked by modifying the feedback graph. The theory posited that, all else being equal, people want

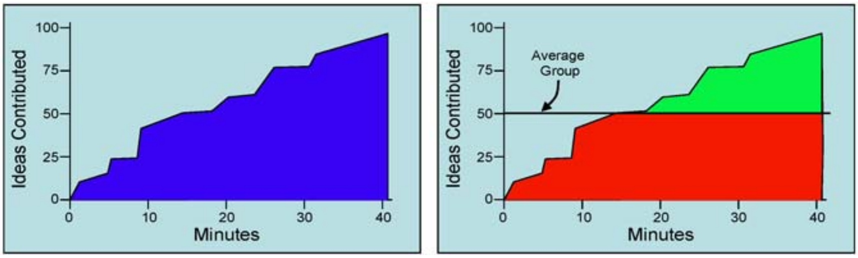

Fig. 3. Left: A feedback graph that provides no basis for social comparison in brainstorming groups. Right: A feedback graph that provides a way to invoke social comparison. Teams that viewed the right-hand graph during brainstorming produced approximately 60% more unique ideas than did teams using the other graph.

the status that accrues from being perceived as contributing fully to a group effort. They want to be seen as stars; they don't want to seem below average. They therefore tend to increase their efforts to at least match the performance of others.

We reasoned that we could add a single horizontal line to the middle of this graph (Figure 3 (Right)), and then tell people, "The average group produces about this many ideas during a brainstorming session…You don't to be below average, do you?" This, we reasoned, should invoke social comparison, causing anonymous brainstorming groups to make more effort, reducing the effects of social loafing. An experiment with 56 groups showed that the groups using the social comparison graph produced about 60% more unique ideas than did the groups using the groups using the standard graph. This was on top of the 50% gain they had already attained by moving from paper to electronic brainstorming. Thus, a good theory led us to a counter-intuitive design choice – the addition of a single horizontal line, which produced a significant improvement in group performance.

3.2 Cognitive Network Model and Creative Problem Solving Techniques

For many years, creativity researchers described creative people, creative environments, creative processes, and creative ideas. This research hinted at, but did not explain what caused creative ideas to emerge in the minds of people working together toward goals. It was not, therefore, possible to predict with confidence whether a new creative problem solving technique or technology might improve creativity. Creativity bordered on a mystical art.

Recently the Cognitive Network Model of Creativity (Santanen, Briggs, and Vreede, 2000) suggested mechanisms of the mind that could give rise to creative ideas. The model drew together standard axioms of cognitive psychology – long-term memory as a web of related concepts , limited working memory, and so on -- to argue that the creative solutions must emerge from novel juxtapositions in working memory of concepts from previously distant parts of the cognitive web. It further argued the number of novel juxtapositions was, among other things, a function of the variety of external stimuli. This theory was consistent with findings that teams using electronic brainstorming technologies could, under certain circumstances, produce substantially more ideas of greater creativity (e.g. Hender, Dean, Rodgers, and Nunamaker, 2002).

The variety of ideas proposed during brainstorming stimulated additional creativity on the part of participants. However, it also suggested that we could make teams more creative if we could find ways to use the collaboration technology that would increase the variety of external stimuli during brainstorming.

In an attempt to use this theory to increase the number of highly-creative ideas a team could produce, we created a new brainstorming approach called Directed Brainstorming. In standard brainstorming, the team responds to a single brainstorming question with a flurry of answers. With Directed Brainstorming, the team receives a stream additional prompts throughout the brainstorming activity. The prompts typically relate to criteria for judging the quality of an idea. For example, for a session on improving factory production quality, the prompts might be:

- *Now give me an idea that would be faster to implement than any you have seen so far.*
- *Now make a suggestion that would be less expensive than those already on the list.*
- *Now think of a concept that would reduce product defects more effectively than any of the ideas we already have.*

A series of experiments showed that, indeed, as the theory suggested, teams using Directed Brainstorming approach produced approximately three times as many unique ideas as those using an un-prompted electronic brainstorming approach, and that among those were a significantly larger number of highly-creative ideas (Santanen, Briggs, and Vreede, 2000). Thus, with a theory to explain our phenomenon of interest, we were able to make non-intuitive choices that carried us well beyond the gains we had already achieved by instinct and experience.

3.3 Technology Transition Model and GSS Architecture

In this case, a small theoretical insight led us to reverse the fundamental assumptions underpinning the architecture of a group support system (GSS), yielding to a new generation of GSS technology. GSS had proven useful in the field, mysteriously, however, despite unequivocal, measurable, profitable results, many GSS facilities fell into disuse within a year or two of being introduced into an organization (Briggs, Vreede and Nunamaker, 2003).

The Technology Acceptance Model (Davis, 1898) posited that people would use a technology to the extent that they found it useful, and that they found it easy to use. The model did not explain, however, why an organization might later discontinue using a technology that still met those criteria. A minor construct of the Technology Transition Model (TTM) (Briggs, et al, 1999) offered an insight that eventually unraveled the mystery. TTM considered ease-of-use in at least three dimensions: perceptual load, access load, and conceptual load. Perceptual load addressed the typical elements of user-friendliness – how easy is it to find and use the features and functions you need. Access load dealt with the fuss-factor – how much effort was required to gain permission and access to the features and functions you need. Conceptual load dealt with understanding what a system is *supposed* to do for the user. Perceptual load for GSS tended to be very low – people with no training could participate in electronic meetings very successfully. Access load was moderate for GSS. People had a bit of fuss getting signed up for electronic meeting rooms and getting the network going.

However, GSS had very high conceptual load. The purpose of the GSS is to create useful patterns of collaboration, yet there was nothing on the screen of a GSS tool telling a user what pattern of collaboration might emerge if a team were to use the tool in a given way. Further, a given tool in a given configuration could be used to create a wide variety of useful patterns of collaboration. Although it took less than a week to learn how to work the system (low perceptual load), it typically took a year of apprenticeship to learn how to wield a GSS in service of group productivity (high conceptual load).

It was a fundamental assumption of GSS design that facilitators would run GSS workshops on behalf of the team, because the facilitators would know how to use the technology to help a group generate, organize, and evaluate ideas, and make informed decisions. However, successful GSS facilitators tend to be bright, articulate, technically-competent people-persons with a penchant for problem-solving. As such, they tended to get promoted away, leaving behind a GSS facility that others did not know how to use effectively.

TTM suggested that if we could find a way to cut the conceptual load for GSS, the technology might achieve wider, sustained use in the workplace. That led us to the concept of collaborative applications built for specific high-value recurring tasks. We reasoned that if we were to create a GSS application that moved practitioners step-by-step through a software requirements negotiation, the users would have almost no conceptual load. They would not have to guess which tool they needed, they would not need to know how those tools should be configured, nor would they need to know which pattern of collaboration might be useful for a given step. A master facilitator could make those choices at design time. At run time, the participants could simply go to work on one fully configured step, and then move to the next when they were ready.

We piloted the packaged-application approach first with paper-and-pencil methods, and then with guidebooks for how to run a particular task on a general-purpose platform. It was clear that this approach cut conceptual load substantially for practitioners of a recurring process. For the first time we observed non-facilitators who successfully conducted their own GSS processes after only a day or two of training, and who subsequently transferred those processes to others in the same organization as the standard way of doing business.

On the strength of these findings, we set about to create a new generation of GSS technology that would support rapid development of completely packaged step-by-step collaboration processes. As of this writing, the new technology has only been released into the workplace for a matter of months. Early results are promising. but it will be years before we know the outcome. It was a good theory, rather than instinct, that led to this radical shift in GSS architecture. If the logic of the theory holds, then we should see more self-sustaining and growing communities of users for purpose-built GSS applications than we did for the facilitator-conducted general-purpose approach.

4 Theoretical Temptations: Models That Do Not Inform

There is nothing so useful as a good theory. A model of cause and effect can suggest ways to design and use our technologies to cause the effects we need. However, all models are not created equal. Our literature is rife with models that yield no useful insight. Such models are seductive, because on the surface, they seem logical. However, in the end, they cannot drive our design choices for collaboration processes and technologies. This section discusses several classes of models that could tempt us into fruitless efforts.

4.1 Grand Theories of Everything

There is a class of theoretical offerings in our literature that propose a plausible umbrella construct, e.g. Group Process, as influencing many phenomena of interest (Figure 4).

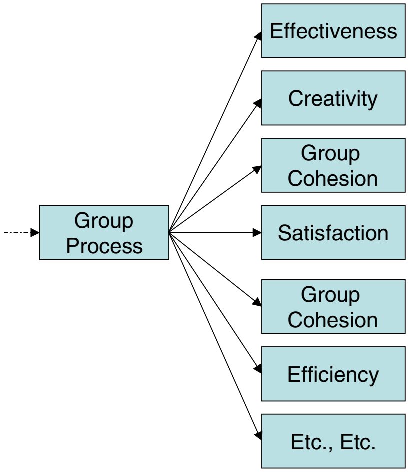

Fig. 4. Grand Theories of Everything posit some plausible umbrella construct as influencing all outcomes. However, such a model suggests no explanation of how to use technology to cause the outcome you need.

It seems logical that group process must, indeed, influence effectiveness, creativity, cohesion, satisfaction. However, the logic of this model breaks down in two ways. Firstly, experience shows that a technology used to improve one outcome will not necessarily improve another. For instance, that which enhances productivity does not necessarily cause creativity (consider the automobile assembly line). And that which produces the highest productivity does not necessarily produce the highest levels of satisfaction (e.g. Connolly, Jessup, and Valacich, 1990). If these outcomes have different causes, then a separate theoretical model is required for each of these outcomes. To produce both outcomes simultaneously, you must understand what causes each, and then seek technical solutions that instantiate the causal mechanisms of both models simultaneously.

Secondly, these models fail the Absurdity Test when one attempts to interpret the box-arrow-box combinations as statements of cause-and-effect, e.g.:

> "The more process a group has, more satisfied the group will be."
> "If we use technology to increase group process, then the group will be more effective, creative, cohesive, satisfied, and efficient…"

A moment of further reflection suggests that at least some of these outcomes may be caused by different mechanisms than others. It may be, for example, that some interventions to enhance efficiency would also enhance creativity, while others would interfere with creativity. Therefore, a single mechanism could not explain both productivity and creativity. A separate theory would be required for each. A model with an outward-fanning peacock-tail of effects should, therefore, at least be subjected to skeptical deliberation before attempting to use it as the basis for a collaborative intervention.

4.2 Grand Theories of Nothing

Other theories in our literature posit general antecedents to a vague abstraction, instead of attempting to explain the causes of a well-defined phenomenon of interest. On the surface, such models may seem plausible at first. However, like the grand theories of everything, these also yield absurd statements when tested, for example:

> "The more knowledge a group has, the more outcomes they will attain."

Models of general antecedents to vague abstractions can be used to discuss group outcomes ranging from brainstorming productivity to nuclear conflict, but they offer no insight about how to cause or prevent any particular outcome. They therefore cannot inform design choices for collaboration technology and process.

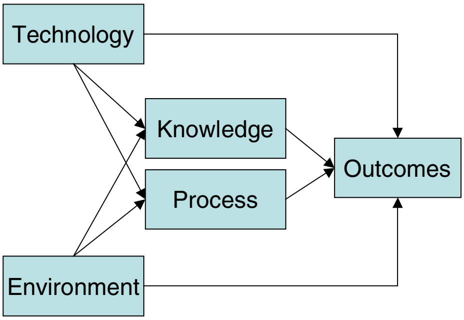

Fig. 5. Grand Theories of Nothing posit general antecedents to a vague abstraction instead of attempting to explain the causes of a well-defined phenomenon-of-interest. Such a model could be applied without fear of repudiation to effects as wide ranging as satisfaction-with-process and all-out armed conflict. However, it cannot articulate the cause of any particular outcome, and therefore cannot inform design choices for collaboration processes or technology.

4.3 Fit and Match Theories

The limitations of the models described in the previous two sections sometimes lead researchers and technologists to yet another blind alley of theory. Experience will reveal the inadequacy of a statement like, "the more technology a group has, the more productive it will become." In practice, sometimes a group with more technology is more productive, but sometimes it is less productive. Those who have the right tools for their task do well, while those who have the wrong task do poorly. It therefore becomes tempting to theorize that the phenomenon of interest (e.g. productivity) must be a function of the match between task and technology. The better the technical capabilities fit the needs of the group, the more successful the group can be. A box-and-arrow diagram of this proposition passes the Absurdity Test that the previous two classes of models fail. The propositions stand up to further thought experiments as a logical, credible assertion.

Match and Fit theories appear in a wide variety of literatures where people consider the use of technology to improve outcomes, and are regularly cited. However, despite their logical consistency, they offer nothing useful to a technology designer. A match theory alludes to patterns of cause-and-effect by implying that an outcome of interest might be caused by one technology but not by another. However, the model stops short of articulating the causes, and sometimes even the effects to which it alludes. It therefore offers the benediction, "Go forth and match," without indicating a basis for creating a match, or for knowing when a match has been achieved.

4.4 Descriptive Attribute, Characteristic, or Factor Theories

There are a number of theories in the literature that posit taxonomies of attributes or characteristics of some object as antecedents to a phenomenon of interest. It is not uncommon, for example to find models arguing that group productivity is influenced by characteristics of the technology, the group, and the environment. There are at least two problems attribute models. First, their propositions do not pass the absurdity test when they are framed as causal statements, for example:

> "The more attributes a group has, the more productive it will become."
> "The more characteristics a technology has, the more productive its users will become."

The logical fallacy of an attribute model is to confuse a category – attributes of the group – with a causal mechanism. Categories are not the same as causes. However, descriptive modes can be useful in a quest for causal mechanism. If research reveals that certain attributes, say, group cohesion and group history, seem to be connected with group productivity, one can ask "Why does cohesion matter? Why does history matter?" Sufficient questioning of this kind can lead to the discovery of underlying patterns of cause-and-effect, making models a useful tool on the quest for good theory.

On their own, however, attribute models suffer another fatal flaw: Infinite decomposability. Any category of attributes can be decomposed into sub-categories. Attributes of the Group, for example, might be decomposed into categories such as group history, group structure, group leadership, and so on. Group structure might be decomposed into emergent patterns, imposed patterns, and other patterns limited only by the creativity of the researcher. As the models become more elaborated with subcategories, they become increasingly difficult to wield, but they gain no additional logical utility as causal explanations, so their usefulness for guiding design choices does not increase.

5 Implications for Collaboration Technology Research

5.1 Theory Should Be Technology-Free

In the early 1990's, some in the collaboration technology reported that GSS tended to make people more productive and satisfied, while others reported that GSS made people more productive but less satisfied, and still others reported that GSS made people less productive and less satisfied (Fjermestad and Hiltz, 1998; 2001). None of these claims were justified. Just as Da Vinci's paintbrush could not cause a masterpiece in the hands of an unskilled painter, so, the outcomes of a GSS session depended on how the technology was used. We had the wrong research question. We were asking, "What is the effect of GSS on group productivity and satisfaction?" Our theoretical models dead-ended until we realized that we had to separate our research question into two questions, one scientific and the other engineering, to whit:

> **Science:** "What causes a group to be productive?"
> **Engineering:** "How can we use technology to invoke those causes of group productivity?"

In the quest for good theory to drive our technology design choices, we must guard against letting our technology creep into both our scientific questions and our theories.

If we ask scientific questions that include technology, we may attempt to build theories that include technology. If we build theories that include technology, then our theories will become obsolete when the technologies become obsolete. Further, our theories will not be able to explain the effects of other technologies, nor to suggest how to change the technology in the theory so as to enable even better outcomes. They will dead-end with the exact technology they embrace. However, if we understand what causes the phenomenon of interest, then we can think about ways to use technology to improve the outcomes we want, and to block the outcomes we do not want.

5.2 Change Research Focus from Technology to Technology-Supported Process

Scientific questions and theories that include technology also lead to unwarranted overgeneralizations that hinder technological progress. Our conclusions early nineties – that GSS caused productivity and satisfaction to increase (or decrease) – overlooked two key points:

a) Any technology that can be used in ways that cause good outcomes can also be used in other ways that cause poor outcomes. (Consider Da Vinci's paintbrush.)
b) An instance of technology is not the same as a class of technologies. Da Vinci cold not accomplish the same effects with a frayed stick that he could with a sable brush, although both could be classified as paintbrushes. In the same fashion, it may not be possible to create the same effects with a one-page brainstorming tool as with a multi-page brainstorming tool, although both could be properly classified as GSS.

Thus, we can never be justified in drawing conclusions about a technology apart from the work-process in which it is embedded. Nor can we defensibly draw conclusions about a class of technologies based on an instance of that class. We can only conclude,

> "When this specific instance of a technology is used in this particular work-process, it produces better (or worse) results than does a different combination of instance and work-process."

It may therefore important for collaboration technology researchers to quickly shift the focus of their research from collaboration technology to technology-supported-collaboration-processes. At this level, collaboration technology researchers can make valid comparisons that produce useful results which can lead to justifiable conclusions.

6 Conclusions

By driving our designs with rigorous theoretical models of cause-and-effect, the field of groupware technology can advance far beyond its already valuable achievements. If we understand the mechanisms that cause our phenomena of interest, we can use a technology in ways to deliberately cause better (or worse) outcomes. If we understand nothing of the causal mechanisms, then we can only achieve a given outcome by accident at first and by rote thereafter. Good theory can make us appear as wizards, able to make otherwise-unexpected design choices that yield better outcomes for our teams. With good theory we may be able to understand other choices don't work out as expected. The key value of the theory driven approach to groupware design derives from the way it discipline our thinking. When we clearly articulate the assumptions and logic that give rise to our technology choices, we will see more clearly than when we do not. We will discover possibilities we never considered. Sometimes it will turn out that our theories are flawed. Sometimes we will make flush out many bad ideas before they consume time, money, and passion.

By embracing technology-driven design, perhaps we can move groupware research from an art toward a science, toward a repeatable engineering practice.

References

1. Agres, Andres; Vreede, G.J.; and Briggs, R.O. "A Tale of Two Cities: Case studies of GSS Transition in Two Organizations." 37th Hawaii International Conference on System Sciences. IEEE. Kona Hawaii, January 5-8, 2004. On CD.

2. Briggs, R.O.: The Focus Theory of Group Productivity and Its Application to the Design and Deployment of Collaboration Technologies. Doctoral Dissertation, University of Arizona, (1994)
3. Briggs, R.O., Vreede, G.J. de, Nunamaker, J.F. Jr.: Collaboration Engineering with ThinkLets to Pursue Sustained Success with Group Support Systems, Journal of Management Information Systems, 19, 4, (2003) 31-63
4. Briggs, R.O., Adkins, M., Mittleman, D., Kruse, J., Miller, S., & Nunamaker, J.F., Jr.: A Technology Transition Model Derived from Field Investigation of GSS use aboard U.S.S. CORONADO. Journal of Management Information Systems. Winter, 1998-99, 15(3), (1999) 151-193
5. Connolly, T., Jessup, L.M., & Valacich, J.S.: Effects of anonymity and evaluative tone on idea generation in computer-mediated groups. Management Science, 36 (6), (1990) 689-703
6. Davis, F. D., Bagozzi, R. P. & Warshaw, P. R.: User acceptance of computer technology: a comparison of two theoretical models. Management Science, 35(8), (1989) 982-1003
7. Diehl, M., & Stroebe, W.: Productivity loss in brainstorming groups: Toward the solution of a riddle. Journal of Personality and Social Psychology, 53, (1987) 497-509
8. Diehl, M., & Stroebe, W.: Productivity loss in idea-generating groups: Tracking down the blocking effect. Journal of Personality and Social Psychology, 61, (1991) 392-403
9. Fjermestad, J., Hiltz, S.R.: An assessment of Group Support Systems experimental research: methodology and results. Journal of Management Information Systems, 15(3), (1999) 7-149
10. Fjermestad, J., Hiltz, S.R.: Group Support Systems: A descriptive evaluation of case and field studies. Journal of Management Information Systems, 17(3), (2001) 112-157
11. Gallupe, R.B., Bastianutti, L.M., and Cooper, W.H.: Unblocking brainstorms, Journal of Applied Psychology, 76(1), (1991) 137-142
12. Gallupe, R.B., Dennis, A.R., Cooper, W.H., Valacich, J.S., Bastianutti, L.M., & Nunamaker, J.F.: Electronic brainstorming and group size. Academy of Management Journal, 35, (1992) 350-369
13. Goethals, G.R., Darley, J.M.: Social Comparison Theory: self-evaluation and group life. In Mullen, B., Goethals, G.R., (eds.), Theories of Group Behavior, New York: Springer-Verlag, (1987) 21-48
14. Harkins, S.G., Jackson, J.M.: The role of evaluation in the elimination of social loafing. Personality and Social Psychology Bulletin, 11, (1985) 457-465
15. Hender, J.M., Dean, D.L., Rodgers, T.L., & Nunamaker, J.F., Jr.: An examination of the impact of stimuli type and GSS structure on creativity: Brainstorming versus non-brainstorming techniques in a GSS environment, Journal of Management Information Systems, 18(4), (2002) 59-85
16. Kerr, N.L, Bruun, S.E.,: Ringlean revisited: Alternative explanations for the social loafing effect. Personality and Social Psychology Bulletin, 7, (1981) 224-231
17. Kerr, N.L, & Bruun, S.E.,: Dispensability of member effort and group motivation losses: Free-rider effects, Personality and Social Psychology Bulletin, 44, (1983) 78-94
18. Post, B. Q.: Building the Business Case for Group Support Technology. Proceedings of the 25th annual Hawaii International Conference on Systems Science, IEEE, (1992) 34-45.
19. Santanen, E.L., Briggs, R.O., Vreede, G.J. de,: The Cognitive Network Model of Creativity: A New Causal Model of Creativity and a New Brainstorming Technique. Proceedings of the 33rd Hawaii International Conference on Systems Sciences. (2000)
20. Valacich, J.S., Dennis, A.R., Connolly, T.: Idea generation in computer-based groups: A new ending to an old story. Organizational Behavior and Human Decision Processes, 57, (1994) 448-467