

HUGH DUBBERLY

MANAGING COMPLEX DESIGN PROJECTS

Editor's note: Last October, Hugh Dubberly made a presentation at the "Living Surfaces" interactive media conference sponsored by the American Center for Design. Dubberly's presentation was on planning and managing interactive, multimedia projects, but his advice is applicable to any design project. With Hugh's permission, we are running the majority of his presentation.

EXPERIENCE HAS TAUGHT ME THAT TWO elements are critical for ensuring the success of any complex design project. First, the people involved must agree on the problem they wish to solve. And second, they must agree on the process for solving it. I want to share with you the tools that I use to facilitate those tasks.

The primacy of the problem

The place to begin any design project is with the problem. Other considerations, such as budget, schedule, technology, skills, media, form, navigation, and even content, grow out of the problem. Everything else is subordinate to the problem. In shorthand: Form follows function.

If you fail to define the problem accurately, no amount of resources—or creativity—will save you. Think of this as the Vietnam phenomenon. During the Vietnam war, Lyndon Johnson put almost unlimited resources at General William Westmoreland's command. But Johnson failed to define the problem he wanted Westmoreland to solve. This management failure had enormous personal cost and almost tore the nation apart.

This failure stands as a lesson for design managers. Well-intentioned, smart people make tragic and costly mistakes when they do not define the problems they wish to solve. The same lesson can be seen a little closer to home. Compare the widespread use of handheld scanner-communicators used by FedEx drivers, with the limited success of PDAs, like Apple's Newton. The FedEx scanner-communicators clearly speed the driver's job and increase accuracy. They meet a need. In contrast, devices like Newton have been unsuccessful because customers don't know what to do with them. Customers don't know what to

do with Newton because the project managers and the designers don't know either. When companies have poor product planning processes, hundreds of millions of dollars are wasted developing devices of which the purpose is unclear. Such disasters are poor management—pure and simple.

The complexity of large projects and the high stakes associated with them create pressures which can easily distract managers from the true problem. The project manager has three primary responsibilities: first, to define the problem accurately; second, to forge consensus on the definition among all the parties involved; and third, to regularly review the work in light of the definition. This role goes by many names: project manager, design

manager, producer, creative director, account executive. Appointing the project manager is an important step in any project. The role does not have to be difficult. A few good tools help make it easy.

Defining the word "problem"

First of all, you need a good model. You need to define the word "problem." God knows, we all have problems. But what, exactly, is a problem? My favorite definition is that a problem is an unmet need.

Needs may be very specific to an individual or organization, but they grow out of broadly shared, basic human needs. In *A Scientific Theory of Culture*, Bronislaw Malinowski posts seven basic human needs: metabolism,

reproduction, comfort, safety, movement, growth and health. Problems arise when one of these needs is unmet. Here we begin to distinguish between two parts of a problem, the need and the means of meeting it.

In *Human Learning*, E. L. Thorndike writes that, "A problem can be said to exist if an organism wants something but the actions necessary to achieve it are not immediately obvious." In *Design Thinking*, Peter Rowe points out that Thorndike's definition gives rise to three general classes of problem: well-defined, or simple; ill-defined, or complex; wicked-hard, or unsolvable.

Well-defined problems are those in which the need is clearly defined at the outset. Only the means of meeting

NEVER PRESENT

CREATIVE WORK TO A

DECISION MAKER

IF YOU HAVEN'T PREVIOUSLY

MET AND DISCUSSED

THE PROBLEM STATEMENT.

the need must be created. Most design school problems fit into this category. Teachers ask students to design a poster announcing an event or to design the interface to an electronic calendar. Both teacher and student make the somewhat dubious assumption that the poster and electronic calendar meet a need. Ill-defined problems are those in which both the means and the need are undefined at the outset. Most professional design projects fit into this category. Unfortunately, most design schools fail to provide courses dealing with ill-defined problems. Wicked-hard problems are a special class of ill-defined problems. Neither the means nor the need is clear, and in addition, the prospect for agreeing on the need is poor. Many social problems fall into this category. Crime, poverty and health-care are examples of wicked-hard problems. In the case of each of these problems, society has not yet agreed on what need is unmet.

During the process of defining a problem, a design manager must ascertain which kind of problem he or she faces. Commercial work on wicked-hard problems carries a high risk and should proceed only with management's understanding the risk. Conversely, well-defined problems should be delegated to less experienced staff. This frees the design manager to focus on ill-defined problems, where he or she can do the most good.

Defining real-world problems

The method that I use to define problems has three parts: collecting input, documenting the need and agreeing on the definition.

Collecting input

Work begins when someone identifies a problem area. Something hurts, so someone asks for help. That leads to an examination, then research to understand the context and extent of the problem. But at this point we only know the symptoms. We still haven't defined the problem.

How do we define the problem? Begin by assembling all the relevant players in a room. If there are too many relevant players to fit in one conference room, hold a series of interviews. Ask each player to describe the unmet need, or in other words, to suggest the cause of the problem. Write down each suggestion. Nothing you will do on the project will be more important.

With each suggestion, ask in turn for its cause. And then the cause of the cause. And then the cause of the cause of the cause. This involves acting like a two-year-old child, in that you must ask "why" after every response. Don't give up easily. Keep at it like a two-year-old. By the time you reach the point where everyone in the room wants to strangle you, you will very likely have found the root cause of the problem.

Total quality management experts often use a fish-bone

diagram to facilitate the process of finding the root cause of a problem. Fish-bones are a good way to structure the responses. It is also a useful mnemonic device for this stage of the process.



Documenting the need

After identifying the root cause of the problem, I find it useful to articulate and document it, in a formal problem statement. I recommend that an individual write the statement and propose it to the rest of the group, because writing by committee can be slow and often leads to verbose and fatuous statements. Ideally the problem is stated in a single sentence—that way it has some hope of being clear and memorable. The form of the problem statement varies depending on the type of problem. I will describe three types of problem statements: cause-and-effect; creative brief and positioning.

Cause-and-effect statements

When working on new or very general types of problems, or problems that seem especially unclear, I resort to a somewhat theoretical but rigorous approach. I review Malinowski's basic needs and organize the problem statement in terms of cause-and-effect. For more familiar or routine problems, I use more specialized forms.

Creative briefs


When working on promotional communications, I organize the problem statement into a slightly longer document called a creative brief. Most advertising agencies use formal creative briefs. Many design firms do not. By the way, you may find it helpful to collect creative briefs from agencies you work with or visit. Some creative briefs are quite long. My favorite is a very short one introduced to me by Wunderman. I remember it as the who-what-what-why brief because it asks those four questions: Who is your audience? What do they believe about your product today? What do you want them to believe about your product after they receive your message? Why will they believe you? When the audience contains identifiable components, the who-what-what-why creative brief becomes a message matrix. Each audience suggests separate whats and whys. Filling in the

matrix is an excellent group activity. It reveals the domain and range of the message and creates a way for all parties to participate and air their views. The message matrix exercise is excellent preparation for writing a positioning statement.

who	what before	what after	why proof
audience 1			
audience 2			
audience 3			

Positioning statements

When working on product communications, I organize the problem statement very much along the lines of classic marketing positioning documents used at Procter and Gamble. I like a positioning statement to follow this form: For [audience], [our product's name] is a [category in which our product competes] that provides [the major benefit of our product] unlike [our major competitor's product].



Another useful tool, the quadrant chart, indicates the position of several products in relation to two axes. Quadrant charts can show dense spaces where competitors cluster and empty spaces where differentiated positions are available.

The value of defining the desired position prior to designing the product is incalculable. Simply put, it means you know what you are doing. If the product is then built so that it delivers on the position, developing communications about the product is very easy. Unfortunately, most companies develop products and then ask the marketing folks to figure out the positioning.

One exception is Honda. In 1990, *The Harvard Business Review* published an article titled, “The Power of Product Integrity.” I highly recommend it to anyone responsible for managing products. The article describes how Honda product managers develop a simple product description prior to designing a new car and how they then use that description to guide the car’s design.

Agreeing on the definition

After you’ve developed the problem statement, you need to be sure to gain consensus on it from all the relevant parties. Failure to get “buy-in” from all the right people at this stage creates the potential for trouble later in the process. Someone who hasn’t agreed on the definition up front is likely to want to change it later.

A primary benefit of these documents comes from using them at creative presentations. The project manager begins the presentation by reminding the group of the problem

statement, creative brief or product position. Then he or she asks the group to confirm the accuracy of the statement.

If the group agrees that the problem statement is still accurate, then the project manager presents the creative work describing how it specifically responds to the requirements of the problem. Likewise the project manager examines any objections to the creative work in light of their relation to the problem statement. Valid objections concern how the creative work responds to the problem statement. Other objections are not valid, unless they lead to reviewing and modifying the problem statement.

If someone in the group objects to the problem statement and the group agrees with the objection, then the presentation should not proceed. The presentation shouldn't proceed because, in such a case, the creative work is no longer relevant. Instead of presenting the creative work, you should focus on redefining the problem. Look at this scenario not as a setback but rather as an opportunity to learn something and to make the creative work stronger. Consultants should make sure their contracts anticipate redefinitions of the problem and allow for appropriate compensation. And, likewise, in-house staffs should make sure decision makers are aware of the process.

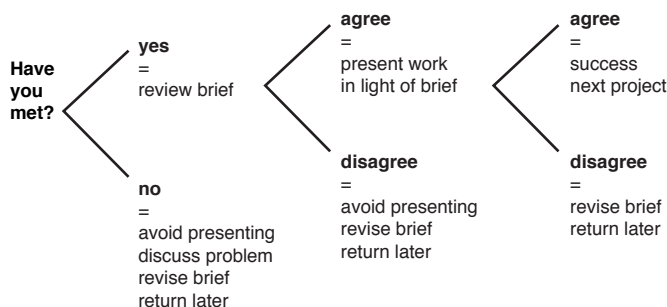
Consider also, this corollary. Never present creative work to a decision maker if you haven’t previously met and discussed the problem statement. If you haven’t met, there is simply no way you can know that you are solving the problem as this decision maker sees it. In this situation, the risk of your presentation being unsuccessful is extremely high. I explain this philosophy at the beginning of a project and use it as a way to ask if the ultimate decision maker is participating in the process. If he or she cannot make time to participate in defining the problem, then something is wrong. It may be that your client is acting without authority or has misunderstood a request. Work should proceed only with the utmost caution.

Process in theory, more models

I’ve devoted so much space to describing how I define problems because I believe it’s the key to the rest of the process. However, the rest of the process also deserves at least as much attention.

To those unfamiliar with it, the design process can seem like a black box. Something mysterious. Something that takes inputs and magically delivers outputs—like a computer.

Some people may see little need to examine the contents of the black box. And there may be little need—if a project is small enough for one person to handle and if that person is experienced with the class of the project. In such situations, intuition and experience will suffice. For some small design projects, idiosyncratic processes may



even be appropriate. But for large or complex projects, and especially for new classes of projects, the process must be explicit or the project will fail.

In design classes, I use a short exercise to begin prying open the black box. I bring to class three office dictionaries, a ream of regular copy paper and a ruler. I give two sheets of paper to each student. Then, as a prelude to the problem, I ask each student to stand, lift the dictionaries and write down an estimate of their weight. We share all the estimates and find the average, which is normally about ten pounds. This prelude involves the students and dramatizes the weight of the dictionaries. Then, I ask the students to use the two sheets of paper to build a structure which will support the dictionaries for at least 30 seconds. The goal is to raise them as high as possible. At first, the students just stare down at the paper. None of them will make eye contact. Some ask about glue, tape, string, rubber bands. The answer is, two sheets of paper only. Finally someone asks about folding or crumpling the paper. The answer, of course, is try it. Once someone tries something, a sort of race is on. They quickly build and test structures. Many of the structures crumple. But failed structures offer lessons. Progress is swift, averaging about 3 1/2 inches every ten minutes, until all three dictionaries are resting 11 inches above the desk supported by nothing more than two sheets of paper.


The students accomplish this rather remarkable feat using two tools that they have learned years before but may not have labeled: the ability to work in teams and a rational design process. Their process, our process, the design process, is at heart trial and error. By describing the process as trial and error, I mean to emphasize its experimental nature. While it is a trial-and-error process, it is by no means random. If we reverse the order placing error before trial, we can make a case that the error-trial model is very similar to a second model, the analysis-synthesis model.

As a general rule the process proceeds from analysis to synthesis. However, this progression is not a two-step sequence. It is not one; two; done. Instead, it's more like one, two, one, two, one, two—a continuing oscillation between two states, an oscillation between analysis and synthesis. An oscillation from error to trial to error to trial to error to trial.

It might also be seen as a loop, where analysis leads to synthesis, which leads back to analysis. This model might be refined into a spiral leading ever closer to a target. The spiral and target suggest that each cycle moves the project closer to a goal.

In their book, *The Universal Traveler*, Koberg and Bagnall expand the two-step analysis-synthesis model by adding a transforming middle step which they call definition. Paraphrasing the authors, "In this step you size up the


situation and develop intentions or guidelines," I would add that you define what needs doing: You define the problem. This is key. It is a pivotal point. You might easily substitute converge, transform and diverge.




The computer business, like most manufacturing businesses, uses an analogous three-step model. The process begins with research, moves into development or engineering, and ends with production or manufacturing

This model works well for engineering managers. Each step may require substantially different skills and thus different project teams. Each step may also represent a separate outlay of capital. Thus the steps can serve as major milestones in contracts. Because each step requires different resources, the project should receive management approval before moving from one step to the next.

While the research and development model has many uses, I believe the analyze-define-synthesize model is more useful for designers. Koberg again expands the model, replacing synthesis with idea generation, idea selection and implementation. Koberg's final refinement is to add a step at the beginning, accepting the assignment, and a step at the end, evaluating the result.



Koberg wisely points out the limitations of viewing this as a linear process. He suggests alternative models that introduce feedback. A loop shows a simple feedback model. A cascade provides a more complex feedback model. Koberg goes on to suggest that "...no stage ever really stops but unlike the [models], each stage is always 'in process' ...the process never ends."




When I was a freshman at the University of Colorado, Cal Briggs taught me a very formal version of this process. Mr. Briggs's process involved preparing an elaborate written document. At the time, I found the process rigid and boring. I simply wanted to design by drawing, and this old guy kept us talking and talking and never let us

draw anything. After two semesters I foolishly left Mr. Briggs's tutelage and moved on to other schools. Although I could run, I couldn't hide.

Fifteen years later, I was trying to explain the problem-solving process to a class of my own when a student, Matt Williams, observed that the problem-solving process is similar to the scientific method. In high school Matt had learned an excellent mnemonic device for the scientific method. I believe Matt's actual words were, "Hey, all you're trying to do is teach us THEOC." THEOC is an acronym for Theory, Hypothesis, Experiment, Observation and Conclusion. Matt had a keen insight. A designer follows a process very much like that of a scientist. Both require experience, rationality and intuition. In a way, the scientific method is simply rigorous common sense. It's the kind of process any good detective follows.

I would like to point out two other, similar processes. First the quality management process made popular by Japanese corporations, Walter Demming and legions of consultants.



The consultant who nabbed me called it the PDCA process. The PDCA process is nothing more or less than design problem solving. Consultants make a lot of money selling this innovation to business managers. Perhaps this is a business opportunity for designers.


A second similar process is the sales process. I had always assumed sales success was essentially a function of personality until Dennis Capovilla, who manages a strategic marketing group at Apple, recently set me straight. Dennis spent several years as a salesman and then as a sales manager. He now teaches a sales management class in the business school at Santa Clara University. Early in the term Dennis buys a new consumer product, say a stereo, and brings it to his class. He describes the product's features to his students and then asks each student to try to sell him the product.

Dennis reports that his students go on at length, some creatively embellishing the product's features. One thing they almost never do, though, is ask questions. They never begin by asking the customer, Dennis, if he has a stereo or if he needs one. After politely demolishing the student pitches, Dennis spends the rest of the semester teaching them a sales process that involves working with customers to define their needs and develop options to fill those needs.

The similarity between Dennis's sales process and the design problem-solving process is astonishing, until you consider how much of what designers do is actually sales.

The basic design process that I generally use looks like this:

gather background information, define problem, suggest options, pick option, create a prototype, test the prototype, repeat steps 1, 2, 3, or 4 as appropriate.



As I noted earlier, the process is less a linear sequence and more a cascade with feedback loops. I believe it also proceeds from vague to concrete, from general to specific. Another way to visualize this idea is as a tree, in which each branch represents a cycle, a cascade, of the process and in which each phase involves more decisions, and more specific decisions, than the last.

A simple model might distinguish between three stages, each involving at least one and probably more passes through the process.

Structure How do we reach the audience?

Content What do we say to them?

Form What does it look like?

A more academic version of this idea is Charles Morris's pragmatic-semantic-syntactic model. Tom Ockerse at Rhode Island School of Design introduces his students to Morris's model to help them understand how signs function. If a sign has these three qualities, then the design process should take them into account. I believe they follow in a natural sequence from general to specific.

Process in practice, an example

We've looked at models for thinking about problems and at tools for defining real problems. We've also looked at models for thinking about the nature of the design process. Now let's look at another model of the design process, one specific enough to be used for managing real projects.

I like to organize the process into three steps, developing the structure, developing the content and developing the form. Getting started is a step in its own right, as is the production and distribution process. The first and last step connect the design process to other issues and other organizations.

The transition from one step to the next is tricky and requires informing everyone involved and confirming

agreement on conclusion of the step. Obviously structure, content and form are deeply intertwined and can only be partially disentangled.

Perhaps it may seem a little silly to suggest a formal process for starting a new project. However, a few years ago, Apple's Creative Services group did not have such a process. Anyone could initiate a project. Managers had difficulty knowing who was doing what and balancing work loads. Much discussion and a few simple tools reduced the confusion. The department created e-mail-based project initiation forms, a central project database, weekly project reports and weekly project status meetings. Project managers were assigned at these meetings and sent to investigate the problem and develop a structure for handling it. They reported back at a later meeting and teams were then assigned.

The process of developing the structure of a project involves initial research to discover the context and domain of the problem; formal definition of the problem; and proposing an overall project process. In this phase, the fish-bone diagram, message matrix, creative brief, positioning document, project proposal, concept diagrams and hardware platform requirements document are all useful tools. Of course it's a good idea to have a review meeting and presentation to confirm the problem statement and project structure before moving to the next phase.

The process of developing the content of a project involves forming a project team; exploring the content domain; documenting, organizing and editing the content; and creating interaction models. Keith Yamashita has shown me the effectiveness of kicking off this phase with a formal briefing for the team and clearly defined brainstorming activities. Keith raises the team's enthusiasm and creates a fun environment with these events. Outlines, flow charts, gantt charts, schedules, user interviews, expert interviews and creative briefs are among the useful tools at this stage.

Some key decisions need to be made here: selection of information models such as sequence, hierarchy, matrix and web; and selection of information types such as text, image, visual narrative, diagram and sound. As the team begins to work out these issues by diagramming the structure of the piece, generous supplies of index cards, Post-it™ notes, and Fome-Cor™ are useful. Already, at this stage, it's not too early to begin considering copyright issues. As before, it's a good idea to have a review meeting and presentation—confirming the organization of the content—before proceeding to the next phase.

The last phase of design, developing the form, involves reviewing the visual elements, creating sample formats, fully defining a visual language and completing a prototype. An audit of existing work, review of existing graphic standards, grids and type specifications are useful tools at

this stage. Central to this stage is user testing and focus group review.

The path from the last design phase to finished piece has its own requirements and pitfalls. I will merely underscore the need for lots of testing and recommend a thorough checklist covering "gotchas" like viruses, time-bombs and file-naming. Just as design is a specialty, so too is production management. My best advice to save yourself some headaches: Hire an expert.

The process I've just reviewed will be useful only if you add to it and customize it. In considering how to customize your process, you may find the following three questions helpful.

First, what kind of problem is it? well-defined? ill-defined? wicked-hard? Next, what factors does the client want to optimize? Paul Pruneau was the first person to introduce me to the canard about time, money and quality. He is rabid in his belief that you can optimize only two of these factors. And that as a result, the third must suffer. Finally, what is the general class of product? Will it be a message, a tool or an environment? Will it be primarily a communication, an interface or a place? These classes yield an entire taxonomy of possibilities, but that's another lecture.

Clearly no single process can stretch to fit more than a fraction of the range of problems that designers face. Likewise, no single method for defining problems is at once specific enough to be a useful tool and general enough to cover more than a few types of problems.

However, the precise steps you use are really less important than writing down the steps you want to use and agreeing on them with your team. The more specific you can be about each step in your process, the more likely you are to avoid forgetting a step, the more likely you are to avoid surprises and the more likely you are to avoid the frustration that comes when team members discover they have different expectations.

The difficulty of finding useful, general models should not deter us from trying to find them. The benefits, in cost savings, in product quality, and in job satisfaction, of articulating our assumptions far outweigh the difficulty of the undertaking. What's more, the new and increasingly complex projects that designers now face, demand familiarity with a range of problem-solving models and clearly articulated processes.

Lighting a candle in the black box of the design process does not destroy the magic. It simply pushes the darkness a little farther away. Of course, it might also ignite an explosion.

