University of Pennsylvania

Penn Libraries
UNIVERSITY of PENNSYLVANIA

University of Pennsylvania
ScholarlyCommons

Management Papers   Wharton Faculty Research

3-8-2006

The Roles of Users, Experts, and Institutions in Design

Karl T. Ulrich
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/mgmt_papers

Part of the Management Sciences and Quantitative Methods Commons

The Roles of Users, Experts, and Institutions in Design

Disciplines
Management Sciences and Quantitative Methods

DESIGN
creation of artifacts in society

Karl T. Ulrich

Contents

Users, Experts, and Institutions in Design

The first act of design was almost certainly *user design*, in that the plan was created by the user rather than by a third-party designer. Perhaps this first user designer contemplated frustration with a task tens of thousands of years ago, formed a plan to address the frustration, and then fashioned an artifact, possibly shaping a stick of wood into a digging implement. A clear distinction between expert designers and user designers emerged at some point possibly first in the domain of architecture. Certainly by the time ancient Egyptians were creating pyramids, the roles of experts and users in design were separated. This separation was probably motivated by the comparative advantage of experts over users in designing enormous structures. The activity of design appears to have become increasingly professional and institutionalized over the next few thousand years. By the 19th Century, as the industrial revolution developed in full, expert designers with specific technical training assumed distinct professional roles, both because of the comparative advantage of expertise and because institutions were formed to exploit the benefits of mass production.

Although a separation between users and designers has increased in many domains over the past several thousand years, the practice of design by users is emerging again in current society in specific domains. This chapter addresses the role of the user in design, with particular emphasis on design by users, and considers how experts and institutions interact with users to deliver artifacts in modern society. My approach is to first lay out a theory of design, based largely on the paradigm of design as *search*. Next, I articulate three modes of engagement by users, experts, and institutions that are exhibited in industrial practice. Then, I outline the drivers of the selection of these modes. Finally, I discuss how emerging technologies and practices are enabling new modes in certain settings, and might enable additional modes in the future.

Design Theory

I adopt an information processing view of design, largely consistent with that articulated by Simon (1996). Within this paradigm, design is part of a problem solving activity beginning with a perception of a gap in the user experience, leading to a plan for a new artifact, and resulting in the production of that artifact (Exhibit MODEL)[1]. I use *artifact* in the broadest sense to describe any product of intentional creation, including physical goods, services, software, graphics, buildings, landscapes, and processes. I include in the model, along with the design process, the production of the designed artifact, as this activity closes the loop between the original gap and the solution.

Exhibit MODEL further decomposes the design process into several elements. This is a codification of a process that may be implicit for many designers, yet these elements can be discerned in some form in most design efforts:

- **Sense Gap.** Design begins with a perception of a gap in the user experience. Without a gap, there is no motive for design. The gap may be perceived by users themselves or by observers.

- **Define Problem.** In effect, problem definition is the creation of a causal model of why the user experiences a gap. This diagnosis can be thought of as an identification of user needs that are not being met in the current state and/or the recognition of criteria for a high-quality solution. Problem definition is implicit in many design efforts, particularly in user design efforts, but is generally an explicit part of professional design efforts, expressed in the form of a design brief, customer needs list, or other document.

- **Search for Solutions.** Given a problem, designers search for *satisficing*[2] solutions. Search itself often includes some form of abstraction/representation. In only a very few domains are search spaces explicitly defined, and in even fewer cases are these spaces finite in scope. For example, the design of internet domain names is con-

---

[1] Terwiesch (2005) provides a comprehensive discussion of product development as problem solving. Product development is a specific economic activity that includes design tasks.
[2] Satisficing is a term coined by Simon (1996) to refer to "good enough" solutions created by agents with bounded rationality.

strained to strings of finite length selected from 36 ASCII characters, an explicit search space of finite scope. However, the design of a custom-built home typically does not face explicit constraints on allowable geometry and may include arbitrary dimensions, and so this search space is infinite in scope. Furthermore, designers of houses rarely work within formal design languages, but rather work with mathematically imprecise representations such as architectural drawings.

- **Select Plan.** Search typically exposes more than one solution alternative and so design requires some sort of evaluation and selection of plans. Some designers consider many alternatives simultaneously when selecting a plan. Others evaluate plans iteratively and select the first plan that satisfices. Sommer and Loch (2004) describe the parallel and iterative modes of problem solving.

Note that in the baseline model, design proceeds from experience to diagnosis to plan to artifact. In modern enterprises, the order is sometimes reversed. The designer begins with an artifact or a plan and searches for needs that the design might meet. This is typical of industries in which effective search methods are lacking, e.g., pharmaceuticals and basic materials. This sequence of problem-solving steps is sometimes called *technology push*.

This design process is typically executed multiple times, as the first artifact produced rarely results in a complete closing of the gap in the user experience. This iteration may occur on a continuum of time scales, ranging from high-frequency iterations by a single individual perhaps over minutes or hours to low-frequency iterations over multiple generations of artifacts within an entire society. For example, Rybczynski (2000) provides a detailed narrative of the evolution of the screw and screwdriver as many iterations of problem solving over hundreds of years.

Design Quality

Design is difficult in that it absorbs substantial cognitive effort, typically requires multiple iterations, and rarely results in an optimal artifact, even in situations for which a formal notion of optimality is possible. The few design domains that have been described by formal representations are, in the nomenclature of computational complexity, *NP-complete* search problems,

meaning that the theoretically optimal solution can not be reliably found[3]. Most design domains have not even been formalized, making the inherent complexity even greater and the prospect of optimality even more distant. However, users can generally still evaluate the quality of the outcome of the design process, and different artifacts designed to address the same gaps can certainly exhibit markedly different levels of quality.
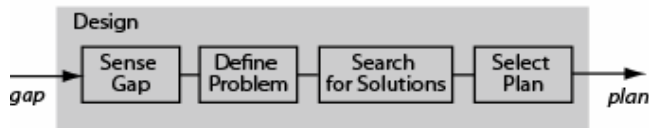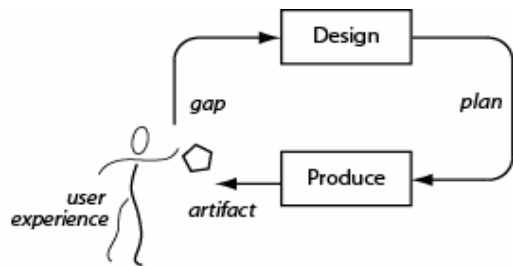
Exhibit Model: Model of design process.

At the most general level, design quality is derived from how well the artifact satisfies user needs, and thereby closes the perceptual gap between a goal state and the current state. The quality of an artifact is linked to at least these characteristics of the design process:

- How well did the designer diagnose the gap in the user experience? Is the problem as understood by the designer consistent with the causes of the gap experienced by the user? In simple terms, did the designer understand the problem?

- Has the search problem been defined in a way that the space of possibilities includes high-quality solutions? In the nomenclature of cogni-

---

[3] *NP* means that the time required for an agent to find a solution increases with the size of the problem according to a relationship that is *not polynomial* (e.g., exponential, factorial, etc.). In other words, the problem "explodes" in magnitude in a way that finding a truly optimal solution is impossible in a reasonable amount of time, even with very fast computing.

tive psychology, has the design problem been *framed* in a way that allows for high-quality solutions to be found?

- Did the designer succeed in finding high-quality designs within the search space that has been defined? Often this result depends on the extent to which a causal model of the relationships between design attributes and user needs can be exploited in navigating the search space. The efficiency of search also depends on the ease and accuracy with which the designer can forecast the quality of a design without actually producing it and having the user experience it.

Although not specifically a risk associated with the design process per se, the fidelity of production of the plan is also a determinant of user satisfaction.

In sum, did the designer understand the problem, frame it in a way that search could potentially find a good solution, find such a solution within the search space, and deliver an artifact consistent with the design.

Another way of thinking about design quality is to identify *defects* that can arise in the design process. For each element of the process, there is at least one potential defect: The designer may fail to accurately diagnose the gap in user experience. The designer may frame the search problem in a way that excludes many high quality designs. The designer may only be able to explore a limited portion of the search space, finding only a few relatively lower-quality solutions. The artifact produced may not be an accurate embodiment of the plan.

Design Modes

I have described the design problem without characterizing the agents that perform the process steps other than referring to them as *designers*. For the purposes of this chapter, I distinguish between *users* and *experts*. Users are the individuals experiencing the perceived gap between the current state and the goal state. They are essentially always a party to the design process[4]. Other terms for users include *customers*, *consumers*, and *stakeholders*, although these terms evoke a more specific commercial context than I intend. Experts have acquired skills and capabilities that allow them to perform most design tasks

---

[4] An exception is perhaps a *design study* done in isolation by a professional designer, but even in this case the designer typically contemplates a virtual user. Design without a user seems to me to be more "my art" than true design.

more efficiently and at a higher level of quality than novices. In some cases an expert may also be a user, but for most design domains this is exceptional.

I make an additional distinction about the institutional context of design. Design may be performed for a specific individual or may be performed for a collection of users. When design is performed for a collection of individuals, some institution is required to coordinate the design and production of the artifact. These institutions are most typically firms, but may also comprise governments, clubs, religious organizations, universities, professional societies, user groups, or even neighborhood associations.

I divide the modes of design into three categories— *user design*, *custom design*, and *common design*— according to the roles played by users, experts, and institutions.

- **User design** comprises a single user designing for his or her own needs. Because the resulting plan is produced for a single individual, and therefore in low quantity, a flexible production process is required to deliver the artifact. Flexible processes need not be technologically intensive (e.g., *flexible manufacturing systems*), but rather need only exhibit relatively low fixed costs for a unique artifact. In many cases, such flexible production processes are craft processes in which skilled people create artifacts with general purpose tools, as is typically the case for unique furniture or unique buildings. An example of a flexible production process enabled by technology is digital printing.

- **Custom design** also comprises flexible production of a unique artifact. However, an expert creates a plan on behalf of a user. In most cases, the user contracts with the expert for this service, as is the case when hiring an architect to design a unique house or a machinery designer to design a unique piece of factory equipment.

- **Common design** differs from custom design and user design in that a single *common* artifact is delivered to a collection of users. Because this common artifact is produced in a relatively large quantity, it may be produced by mass production methods, processes which typically incur substantial fixed costs for each variant of the product, but relatively low marginal costs of producing additional units. Common design involves an institution of some kind, usually a firm, that assesses

the gaps in a set of users, creates a common plan for addressing those gaps and delivers a common artifact to those users.

This taxonomy focuses on differences in the way design is performed, and I do not distinguish between flexible production by users and flexible production by experts. Mass production because of its very nature must be performed by an institution of some kind as it serves a collection of users with a common artifact.

These categories are intended to be exhaustive and mutually exclusive relative to the variables identified here. However, all three modes may exist to serve different individuals within the same community of users or market.
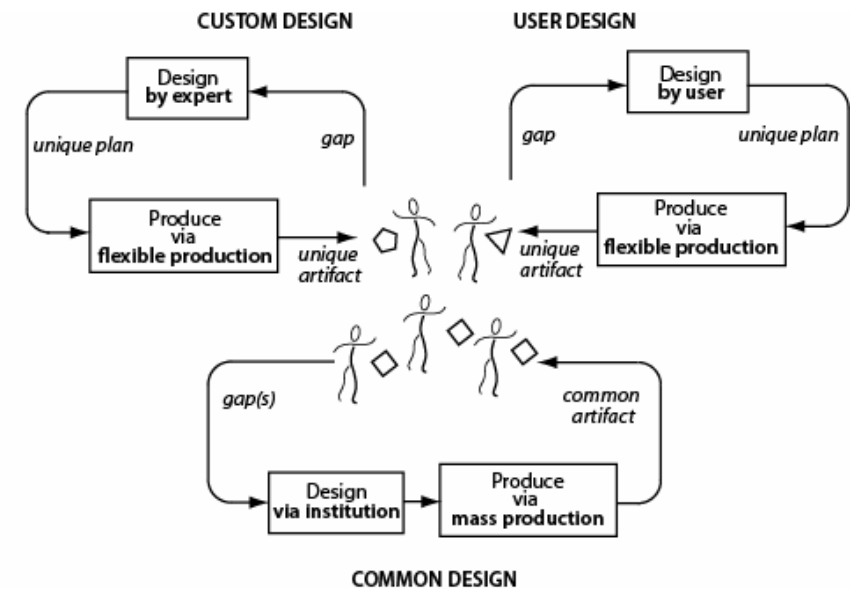
Exhibit Modes: Three modes of design which may be exhibited within a community of users.

Drivers of Mode Choice

Assuming that historically the first design was user design, why did the other modes evolve and why do they exist? What are their relative advantages? What drives the choice of mode in a particular setting?

Economies of scale in production lead to institutional design

A very large fraction of the economic value in retail trade in current society is through just a few very large distribution channels (e.g., Walmart, Target, Home Depot, Carrefour). Most products in these channels are produced in high volume (e.g., 10k to 10M units/year) for a mass market. This is because for these products mass production offers a crushing advantage in satisfying user needs at low cost. This advantage arises because of economies of scale in design and production. Creating 10,000 pairs of identical shoes can be 100 times less expensive on a per-unit basis than creating only 1 pair of unique shoes. Very few consumers have distinct enough needs to be willing to pay a hundredfold premium for shoes made uniquely for them. In sum, the cost structure of most design and production processes provides a compelling motive for clustering similar groups of users and addressing their needs with a common design.

A common design requires an institution of some kind, because to achieve commonality, users must be grouped, the gaps in their experiences assessed, and a common artifact designed and produced for them. In sum, economies of scale lead to mass production; mass production requires a common design; a common design requires an institution. For this mode, user design is not generally possible. To the extent that design is performed by a single individual, or even by a team, the remaining individuals whose needs are addressed by the common artifact will not be designers. Instead their experience will be assessed vicariously by others in the common design mode.

Advantage of expertise in design drives the selection of the custom design mode

Design is performed for a single user when that user's needs are unique enough, given likely economies of scale in design and production, that a unique artifact is preferred to a common artifact (Lancaster 1990). This case arises frequently in architecture (custom homes, buildings, landscapes), food, software, and graphics. This mode is also exhibited occasionally in furniture, apparel, sporting goods, and tools. It is exhibited rarely in home appliances, automobiles, aircraft, medical devices, or computers, domains for which the economies of scale present nearly insurmountable barriers to

unique artifacts, even for the very wealthy[5]. The design of a unique artifact in this context may be performed either by the user or by an expert on behalf of that user, leading to the two modes in the upper half of Exhibit MODEL.

All other things equal, design professionals develop expertise that allows them to perform design tasks better than novices (Ericsson 1996). Given that most users will be novices, experts will outperform novices in most design tasks. However, costs are incurred in engaging an expert, and so the expert design mode will only be selected when the advantages of expertise outweigh the costs of engaging the expert. These costs can be thought of as *direct costs* paid to the expert and as *transaction costs* associated with retaining the expert. Direct costs are straightforward. Most experts will be paid for their services. A "do it yourself" (*DIY*) user values his or her required design effort at less than the cost of retaining the expert, accounting for possible differences in the resulting design quality.

Transaction costs are more subtle. Transaction costs are incurred in defining a design problem and in evaluating alternative solutions. On first reflection, a user would appear to have an advantage over an expert in diagnosing the gap in his or her own experience. I believe that this is sometimes true but not necessarily so. Experts by definition have encountered similar design problems many times before and will likely have observed empirical regularities in user needs. Experts typically also deploy techniques for probing user needs, such as interviews and observational methods (Ulrich and Eppinger 2004). In many cases user needs are *latent*, in that they can not be spontaneously articulated by users, but if these needs were satisfied the gap in the user experience would be addressed. Of course, a risk of expertise is that it frames the designer's diagnosis of the problem. An architect may define a gap in the communication patterns within an R&D organization as a problem relating to the built environment, whereas a management consultant may define the same gap as a problem of organizational structure.

Search almost never results in a single plan, but rather exposes several alternatives which are promising enough for serious consideration. Evaluation of alternatives typically occurs "on paper" before an artifact is produced. Once an artifact has been produced, there is almost always an

---

[5] Some artifacts can be decomposed into a platform and derivatives, with the platform a common artifact and the derivative a unique artifact. In a subsequent section, we discuss hybrid modes of design, which can arise in such cases.

evaluation through testing by the user. Users are clearly best at assessing, through their own experience, whether an artifact actually closes the sensed gap in their experience. While experts may productively observe patterns in behavior, ultimately the user is the frame of reference for the gap in the first place, and is the only agent who can conclude that the gap has been addressed. However, users are typically ill equipped to forecast the extent to which a design alternative, represented abstractly, will meet their needs. Because they do not work daily with design representations, most users are not skilled at visualizing an artifact, at mental simulation of the artifact's function, and are not alert for common pitfalls for a category of artifacts.

Given these characteristics of transaction costs, users are actually likely to have an advantage over experts when design alternatives can be readily generated and when plans can be accurately evaluated quickly and at low cost, as when realistic prototypes can be produced readily. In such environments, the user can achieve high-quality design through rapid iteration and learning. Expert design in the same context can incur high transaction costs because of the switching back and forth between search by the expert and evaluation by the user. In this situation, the more efficient search by an expert may be outweighed by the reduced transaction costs of user design.

An additional driver of user design is the utility (or disutility) some users derive from solving their own problems. To the extent that there is a psychological benefit derived from the process of design ("I designed it myself!"), a user may be willing to accept a lower quality outcome even at the same cost of expert design.

Synergies among modes

All three modes of design can and typically do exist in the same community and for the same category of artifacts. Some people engage in user design. Some people engage in custom design. Everyone participates in common design, at least through their consumption and use of artifacts.

A commonly occurring pattern of innovation is for a new artifact to emerge through user design and then to be adopted, often with some refinement, as part of a common design effort. This process of appropriation and improvement may take place over many years and even generations. This pattern of innovation has been documented in detail by von Hippel (1988). However, the migration from a unique design to a common design need not originate in user design. An essentially similar pattern involves the migration

from expert design of a unique artifact for a single user to common design by an institution for a collection of users. In either case, an individual user uncovers a set of user needs and a design that addresses those needs. This design is subsequently exploited by an institution to deliver a common artifact.

Hybrid modes

An artifact may be the result of more than one mode of design if it is comprised of more than one element. For example, a common component may be used in combination with a custom component. Or, one or more attributes of a component may be customized, with the rest standardized. This approach is sometimes called a *platform strategy* and is closely related to the notion of *mass customization*. By adopting this strategy, a producer may be able to offer a user a unique design while exploiting the economies of scale associated with the standard elements of the product. Randall, Terwiesch, and Ulrich (2005) provide a detailed discussion of user design for customized products.

Enabling Processes and Technologies

Mode choice in design is strongly influenced by changes in design and production processes and technologies. New technologies and processes have emerged in the past few decades that are changing the way design modes are adopted in practice.

Templates

The problem of search is dramatically simplified if a *template* is adopted. A template is a fixed architecture for an artifact within which alternative elements may be placed (Ulrich 1995). For example, iPrint is a web-based system by which users may design printed items such as business cards, stationery, and party invitations (Exhibit IPRINT). Each of several types of items is represented with a standard template. Within that template, choices may be made of typeface, type size, colors, position of graphic elements, paper, and textual content. By constraining search to a selection of elements within a fixed template, the design problem is bounded sufficiently that many users find that they are able to find satisficing solutions without retaining an expert. Digital printing technology is sufficiently flexible that unique artifacts may be produced in relatively low volume (50-1000 units) at reasonable cost.

Exhibit IPRINT: Web-based interface for user design of a New Year's party invitation based on a template.

Design Grammars

A design grammar is a set of rules defining "valid" designs, including a definition of the elements of the design and the rules by which they may be configured. (A template is a very restrictive type of grammar in which the alternative selections of elements must always be configured in the same way.) Grammars have been developed and used for VLSI circuit design, for computer system design, and for chemical process design. Formal grammars have otherwise rarely been used in design practice. However, the develop-

ment and use of such grammars offers the prospect of making search more tractable for novices, or even computers[6].

Stiny (1978) developed a design grammar for several domains in architecture, including Queen Ann style houses. Exhibit STINY is an example of several instances of valid Queen-Anne houses within Stiny's grammar, each showing a different valid porch configuration for a single main house plan.
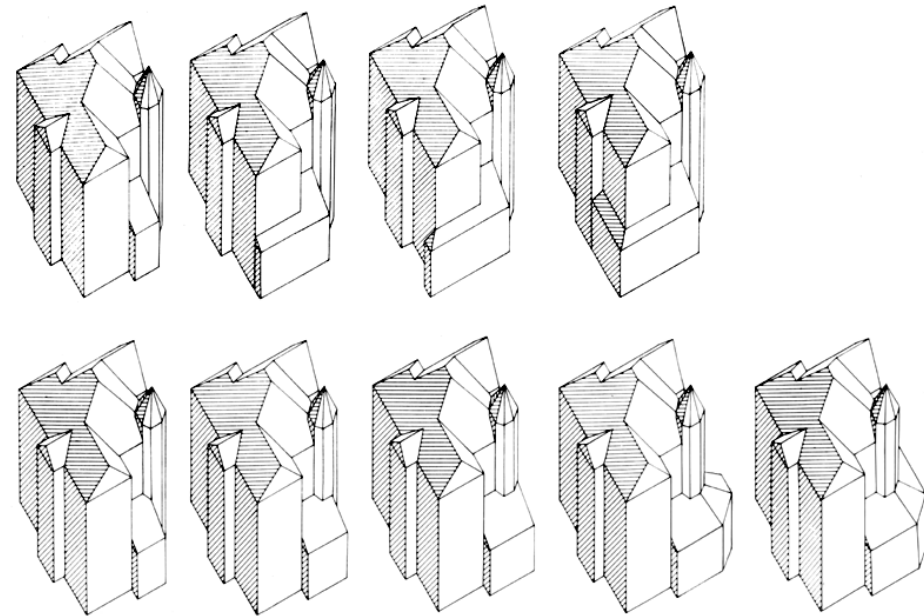
Exhibit STINY: A few instances of a "Queen Anne" design composed within the Queen-Anne grammar. Source: Stiny (1978).

A grammar defines a universe of valid designs. While it may enable efficient search, it also restricts the space of possibilities to the scope of the grammar. Consider the designs of Frank Gehry such as the MIT Stata Center (Exhibit STATA). In the late 20th Century, Gehry's work appeared fresh precisely because it deviated from existing grammars, possibly the way the Queen Anne style appeared fresh in the late 19th Century. Interestingly, over his career Gehry has designed enough buildings that one can start to imagine a formal grammar defining a valid "Gehry style."

---

6 Goldenberg and Mazursky (1999) make a compelling argument that what they call "templates" (actually closer to a grammar in my nomenclature) can be used to characterize successful designs for advertisements and new product concepts.

Exhibit STATA: The Stata Center at MIT, designed by Frank Gehry. Source: http://yoda.zoy.org/copynotice.

Search Automation

If a design domain can be formalized through a design grammar, then the prospect of automating search emerges. A second requirement for automating search is that a formal evaluation function (or *objective function*, in the language of optimization) can be articulated. Without some way of automatically estimating the quality of a design, automating search is unlikely. For highly structured design problems, such as creating a customized personal computer to meet the needs of an individual, search automation is currently feasible (Randall, Terwiesch, and Ulrich 2005). Additional problems are likely to be addressed by search automation in the future.

Rapid Prototyping

Most design efforts require the designer to forecast the extent to which a contemplated alternative will satisfy the needs of the target user. A forecast is required when the cost of producing the artifact, even in prototype form, is relatively high. Rapid prototyping technologies, which might be called more appropriately *inexpensive* prototyping technologies, allow the designer to produce relatively more prototypes for actual testing and can therefore reduce the importance of accurate forecasting of design quality. In the hands of a novice designer, the act of testing many prototypes can substitute to

some extent for expertise in search and evaluation of designs and thereby enable user design where custom design or common design was previously the norm.

Exhibit SLS shows several chess pieces made directly from computer models using the *selective laser sintering* (SLS) process. The cost and time required to produce physical models of complex geometric forms like these has fallen by at least a factor of ten relative to conventional prototyping technologies (in this case, carving by hand), enabling more frequent evaluation of physical prototypes as opposed to requiring the designer to completely refine the form of an object before committing to an expensive and time-consuming prototyping process.


Exhibit SLS: Chess pieces fabricated using the Selective Laser Sintering (SLS) process, a rapid prototyping technology. Source: http://www.kinzoku.co.jp/image/zoukei_p3_b.jpg

Flexible Production

Flexible production is a means of producing artifacts with relatively low fixed costs per variant of the artifact. For example, laser printing of documents is quite flexible, allowing 10 different documents to be printed at about the same cost as 10 copies of the same document. Computer-controlled laser cutting machines allow arbitrary trajectories to be cut in plywood, sheet metal, and plastic sheet, with essentially no *set-up cost*. To the extent than an artifact can be produced by flexible production means, unique artifacts can be produced for individual users at reasonable cost. Flexible production technologies therefore enable custom design and user design. Exhibit CNC shows a web-based design interface that creates instructions for a computer-controlled milling machine, which can be used to flexibly produce three-dimensional shapes as shown. CNC milling is a material removal process incurring only modest fixed costs per variant of the artifact and therefore enabling relatively low volume production.

Tournaments

Tournaments in design have increased in popularity with the advent of mass media channels, but have probably been used by institutions for a long time. In a tournament, many individuals or teams submit plans or prototypes, which are typically evaluated by experts, sometimes with panels of users, and sometimes through testing. Some tournaments are intended to be primarily design mechanisms for a producer or user. Examples of these competitions are QVC's product road show, which visits 10 cities in the United States each year to screen new products, and the U.S. government agency DARPA's *Grand Challenge* autonomous robotic vehicle competition. Other tournaments are intended primarily to deliver entertainment to an audience. An example of this type of competition is *Million Dollar Idea*, a televised competition in which a winner is granted $1 million to commercialize his or her invention. Tournaments exploit large numbers of parallel searches by individuals, sometimes collecting design alternatives from thousands of entrants. This strategy can be particularly powerful when seeking new ideas for products in that a raw plan, perhaps only partially developed, can be selected from the efforts of many individuals and then refined professionally in through common design by an institution. Tournaments may also exploit a tendency by entrants to overestimate the probability of success, possibly resulting in more design effort per unit of investment by the tournament sponsor than could be achieved by other means.
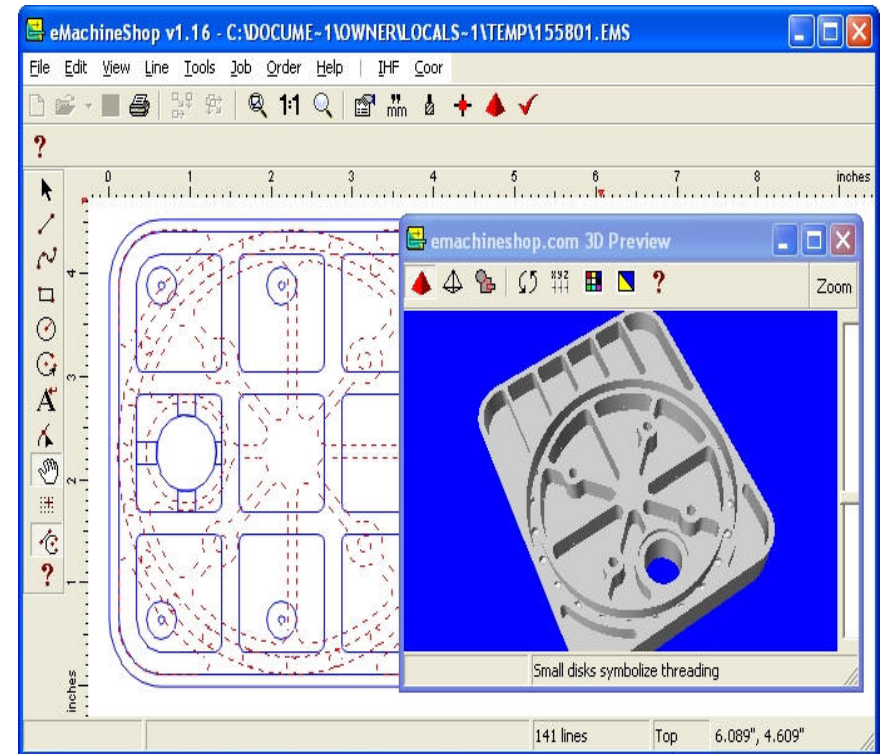
Exhibit CNC. Aluminum part flexibly produced by a CNC milling machine. A web-based design program can be used to create instructions for the milling machine. Source: emachineshop.com.

Open Source

The practice of *open source* arose in the software engineering community and comprises, at a minimum, the free publication of the "source code" for an artifact. For software, the source code is the program instructions in human-readable form, typically as they were written by the designer. For documents, the source code is the text, in readable, editable form. For a physical good, the source code might include geometric information, materials specifications, control algorithms, and/or process specifications.

The rationale for open source is that some users will sense opportunities for improvement in an artifact and will themselves make those improvements. Several open-source communities have developed and are active, with the most famous being the Linux computer operating system. Most of these communities have some mechanism for evaluating and ratifying potential improvements submitted by members of the user community. Remarkably, some open-source artifacts evolve with almost no managerial oversight. For example, the Wikipedia encyclopedia is open source, and can be modified by anyone in the world with access to an internet browser. Open source communities need not be firms, but they are nevertheless institutions that enable the common design mode.

Design Kits

Design kits are tools to facilitate the design process, often provided at no charge by firms seeking to produce the unique artifacts of designers, or who otherwise benefit from active design communities. Producers of specialized semiconductor devices will sometimes provide designers with "breadboard" systems incorporating the devices to enable experimentation and trial, and in the hopes that these devices will be used in a new artifact. Design kits reduce the fixed costs of designing a unique artifact and so enable expert design and user design.

User Groups

User groups are sets of users with communication mechanisms to facilitate the exchange of information relative to a class of artifacts. These mechanisms are increasingly electronic, typically implemented via the internet. User groups are often structured around issues or questions, sometimes called discussion threads, although some user groups have formal administrative elements such as managers and committees. User groups enable user design by allowing plans from one user to be communicated to another with similar needs. User groups can also facilitate common design by allowing users to share information about gaps, coordinate plans, and even test prototypes.

An example of a user community is flashkit.com, a community of designers using the Macromedia Flash multimedia programming language. As of this writing there were about 500,000 members of this community. In this case, a primary beneficiary of the user group is the firm Macromedia.

Concluding Remarks

This chapter articulates the modes of design adopted by users, experts, and institutions in creating new artifacts. User design is a tantalizing prospect by which users create unique artifacts to address their own needs. Yet, expert design and common design remain prevalent modes. The choice of a particular mode is driven by the comparative advantage of experts, by economies of scale in design and production, and by the transaction costs of engaging experts, features which remain the foundations of modern economic life. However, emergent processes and technologies such as rapid prototyping and design grammars can alter the economics of mode choice.

References

Ericsson, K. A., The Acquisition of Expert Performance: An Introduction to Some of the Issues. In K. A. Ericsson (ed.), *The Road to Excellence*, Lawrence Erlbaum, Mahwah, NJ, p. 1-50, 1996.

Goldenberg J., Mazursky D., Solomon S., "Creative Sparks," *Science*, Vol. 285, Issue 5433, September 1999, p. 1495-1496.

Lancaster, Kelvin, "The Economics of Product Variety: A Survey," *Marketing Science*, Vol. 9, No. 3, Summer 1990, p. 189-206.

Randall, T., C. Terwiesch, and K. Ulrich, "User Design of Customized Products," *Marketing Science*, forthcoming 2005.

Rybczynski, Witold, *One Good Turn: A Natural History of the Screwdriver and the Screw*, Scribner, New York, 2000.

Shape Grammar Website, http://www.shapegrammar.org/index.html, accessed October 10, 2005.

Simon, H. A., *The Sciences of the Artificial*, Third Edition, MIT Press, Cambridge, MA, 1996.

Sommer, S.C., and C.H. Loch, "Selectionism and Learning in Projects with Complexity and Unforseeable Uncertainty," *Management Science*, Vol. 50, No. 10, 2004.

Stiny, G. and W.J. Mitchell, "The Palladian Grammar," *Environment and Planning B,* Vol. 5, p. 5 – 18, 1978.

Terwiesch, C., "Product Development as a Problem-solving Process," in S. Shane (editor), *Blackwell Handbook on Technology and Innovation Management*, forthcoming.

Ulrich, K.T. and S.D. Eppinger, *Product Design and Development*, Third Edition, McGraw-Hill, New York, 2004.

Ulrich, Karl, "The Role of Product Architecture in the Manufacturing Firm," *Research Policy*, Vol. 24, p. 419-440, 1995.

Von Hippel, Eric, *The Sources of Innovation*, Oxford University Press, New York, 1988.