The semiotics of umbrellas

john wilkes

HP Laboratories, Palo Alto, CA, February 2003

It's always more fun to tilt at an appropriately-sized windmill – and agreeing on which windmill to tilt at often makes the difference between success and failure in research. What I offer here is a humble suggestion for some vocabulary with which to discuss windmill tilting, in the hope that the endeavor will be more productive for all concerned if the beast can better be identified, named, and communicated about.

Once upon a time, I found myself engaged in a discussion with a colleague on the relative merits of prototyping a piece of software. That conversation proved unfruitful: as we later discovered, we had very different ideas about what was meant by the term "prototype". One of us was convinced that the only prototype of value would have to be a first-cut of the software that could be shipped as a product, after some engineering had been "applied" to it. The other was equally adamant that prototypes were merely vehicles to get across an idea – a way to sell a proposal, and perhaps either to demonstrate that it did something useful, or to determine if it did. We parted company, each mystified at the others' intransigence.

Some time later, having become older, if not wiser, I, and a new team of people that I was privileged to work with, decided that this was all simply a confusion over vocabulary, and that banning the "p-word" would serve us all well. Indeed it has – but we never really found a satisfactory replacement for it that we could remember from one day to the next.

Then, a couple of years ago, something clicked after one of those interminable discussions about "what should we do next". The images shown here emerged (a little soggy) the following morning during my shower. Just for fun, I've reproduced my first scribbles of them – illegibilities and all, together with their definitions, and a few related thoughts.

Research nugget: a coherent unit of research work, and typically the result of a small(ish) research project – maybe as little as a person-month's work, maybe as much as a dozen or so. Often achievable by a person or three; good ones can result in nice technical papers.

Connotations of gold mining are not unintentional.

Testbed: a vehicle for obtaining research results as rapidly and efficiently as possible. The purpose of a testbed is to develop, nurture, and support one or more research nuggets – nothing more. Although a testbed may be (too often is) pressed into service in other roles, such as showing off the research work, this mixing of purposes is best viewed for what it is – a distraction. The evaluation criterion for a testbed is the ease with which research can be performed. (To help get this across, in a UNIX-centric culture, we used to say "if MS-DOS works better, use it"!)

Vision: a description of some goal, a result that a project is trying to achieve – an "end state" in the consultant's jargon. I've found it helpful to separate out the vision from the research – the research (at least in my world) is best thought of as supporting or enabling a vision. (Indeed, it often comes about by working backwards from a vision – "what's impossible in that vision, today?")

Visions are helpful in justifying work – explaining "what it all means" and why we want to go there. Good visions seem to be contentious and attractive – bad ones vacuous, or simply dull. Visions are good vehicles for teasing out subjective notions of "value" from possible participants in, or customers of, a piece of work: if a vision doesn't catch people's imagination, the work to achieve it is unlikely to be pursued with enthusiasm.

Multiple smaller visions can cover a bigger one
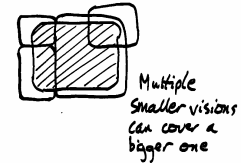
It's usually helpful if there is a common vision, since that means people who subscribe to it agree on the goal. But associated with the one larger vision, it's also common to have multiple, smaller-scale or smaller-scope visions. Ideally(!) the smaller visions complement one another, and can be seen as contributing to the bigger one.

showcase ("demo")

nuggets fit into (one or more) showcases

Ideally, showcases cover the vision
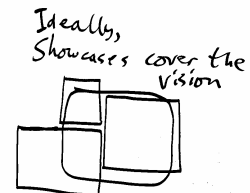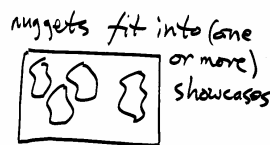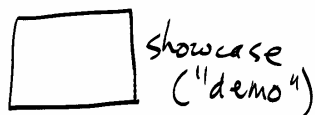
Showcase: a tool used to demonstrate some or all of: (1) a vision, (2) research work, and/or (3) that a team is making progress. A showcase that's an executable piece of code is sometimes called a demonstrator. Other forms include published papers, mocked-up user interfaces, storyboards, and presentations (preferably with attractive animations) claiming magnificent things. The test for a good showcase is that it makes visible what the excitement is all about, and focuses attention on the accomplishments, rather than on the effort required to achieve them. Unfortunately, this seems to mean that it's quite hard to build good showcases for operating systems, middleware, or anything that hides or reduces work.

In some cases, a testbed may be usable in – or even as – a showcase. But these two roles are different, and suggestions to "economize" in this form should usually be treated with skepticism: it's all too easy to end up with an unconvincing showcase that it is inconvenient to do research in.

Most research nuggets can fruitfully fit into one or more showcases. Indeed, it's often a useful idea to think through how the research will be demonstrated before too much effort is put into doing it!

Showcases can readily complement one another: a larger vision may best be described and demonstrated in pieces, especially early on, although it's often helpful if there's a "core" showcase that is being aimed for, and some of my colleagues have reported that mocking up such a showcase is often all that it takes to sell a key idea.

Umbrella projects: a grouping or coalescing "wrapper" that ties together a set of other activities into a common theme. Like the p-word, the "umbrella" concept often seems to

cause confusion. Indeed, I've heard it used to describe a vision, a single large project, and a politically-correct shield for continuing business as usual (especially after inputs of the form "it is now a corporate mandate that all projects must …")! More useful, perhaps, are the relatively benign forms described here.

1. Flying in formation: here, there is a set of research nuggets that share a common vision, but the ties between the pieces of work that go into the nuggets are relatively weak, and it's unlikely that a single, coherent showcase is put together.

Compared to the next alternative, the lack of a single showcase can greatly reduce the amount of integration work required, but it still may be possible to spin (sorry: present) all the research as conforming to a single coherent vision.

2. A unifying showcase: this is closer to the "single large project" model – a single showcase is used to tie together the individual pieces of research, and demonstrate them and their inter-relationships.
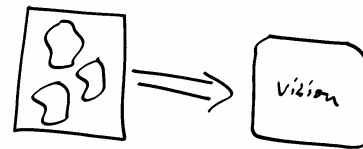
This is usually significantly more work to get set up – especially for the first few nuggets – but can present a correspondingly more compelling façade. In my experience, getting one of these unifying showcases agreed to is a black art. It requires somebody to have the courage of their convictions – and a silver tongue – to persuade others of the viability, utility, and excitement of the associated vision. It can be done. I wish it were done more often. The (slightly) greater ability to pull this off is one of the few distinguishing factors associated with a top-notch industrial research establishment, as compared to an academic one.

If effective, such a unifying showcase has the advantage of achieving higher impact than a single research nugget can manage by itself. The obvious disadvantages are the relatively high risk ("what if we pick the wrong problem?"), exacerbated by the fear of putting too many eggs in one basket; the difficulty of reaching a common understanding of the goal ("what about this other interesting side issue?"); and the potentially high integration cost of the showcase artifact, which now becomes more of an industrial-strength vehicle than a research tool *per se*.
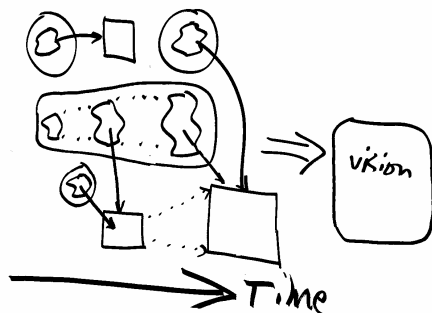
In practice, of course, nothing is as simple as this exposition suggests, as the (deliberately rather muddled) diagram to the left attempts to illustrate. Real-life projects mix and match approaches and techniques, in response to all sorts of outside and internal pressures, requests, and ideas. Research nuggets, testbeds, and showcases come and go – or morph into new ones as understanding, interest, and opportunity allow.

Most projects end up with a mixed bag of assorted testbeds, supporting changing research nuggets that contribute to different showcases at different times. But good projects seem to retain at least a thread of a common vision – even if parts of it may be submerged temporarily, and new elements appear.

I've found that attempting to tease out the different roles and assumptions of each piece is still a beneficial activity. Recursion is often useful in this exercise: what looks to be a research nugget (or vision, or …) can often be sub-divided, and the same analysis applied to each piece.

Over the past couple of years, these ideas have seemed to resonate with my colleagues, and they have proven useful as a way to communicate ideas for structuring and focusing some of our work. One day, perhaps, they might help us approach the scale of effects and impact achieved by the apparently effortless, *laissez faire* project management processes that the Computer lab used in the heyday of the Cambridge Distributed System. We can but dream.