

Back to the Future

A Journey from Science to Craft . . . and Back?

Future Internet Assembly

John Day

Budapest Hungary

May 2011

Any sufficiently advanced craft is
indistinguishable from science.

- after Arthur C. Clarke

The Professionalization of the Sciences

- Lee Smolin points out in **The Trouble with Physics** how professionalization has led to selecting for “master craftsmen” to the almost complete exclusion of “seers.”
 - Master craftsmen are good at math and science; solve problems quickly.
 - Seers are also good at math and science but are more interested in getting to the bottom of things, what is really going on.
 - Our technology has become so advanced that we mistake craft for science.
- Combined with groupthink, its effects are an alarming trend.
 - Everyone is seeing it.
- In computing, this has translated to selecting for fast (but not necessarily good) coders that can find a solution to a problem quickly.
 - Anything that works (and is deployable in the Net) is acceptable.
 - Whether it furthers understanding or lays a foundation is not relevant.
 - Incremental change and running code selects for craftsman.

Our Field Has Become a Craft

- We don't select for those who are good at finding solutions that lay the groundwork for the next step, and want to know why these are the solutions, the patterns in them, what is the range of solutions and when are they useful, i.e. theory.
 - Our idea of moving toward simplicity is a simpler patch
 - This process has been aided by Moore's Law, which has allowed us to ignore fundamental questions.
- This often leads to an answer but not to understanding the answer.
 - A reticence to vary from the one answer we know (sometimes called tradition).
 - But getting product out the door doesn't allow for such niceties.
 - No, it doesn't.
 - that is why it isn't research and certainly isn't academic research
- We have a craft.

The Role of Theory is Unique and Critical to Science

- Proof and by implication theory are unique to Western civilization.
 - Other civilizations developed sophisticated technologies and mathematics, but as recipes, steps to be followed. They were artisan traditions.
 - Thank goodness the Greeks were an argumentative bunch!
 - Theory provides structure and cohesion to knowledge.
- Theory
 - Reduces what must be remembered
 - Reveals what we know and don't know
 - Allows insights that were obscured without it
 - Most importantly, theory makes predictions.
 - And raises more questions
- Science provides the framework for engineering.
 - Engineering without science is craft.
- If there *is* an arrow of progress, theory is its vector field.
- Because without it,

There is Stagnation

- History shows, craft tends toward stagnation, relying more on tradition than analytical inquiry, e.g. China pre-Western contact, (Needham)
 - The West was only one to make the leap from artisan to science
 - Proliferation of technique leads to lost knowledge and stagnation.
 - We see evidence of this in our own field. In fact,
- The failure of the Future Internet efforts is another example.
 - Failure?
 - A decade of work has not yielded a single breakthrough
 - Remember all the wondrous predictions of radical new models, clean-slate, fresh thinking?
 - With nothing to show for 100s of millions spent, now we hear:
 - “The Internet has always done best by evolving. There isn’t that much wrong with it. Fixing a few things here and there and all will be okay. We can evolve out of the current problems.”
 - This is CYA,* not research.
 - Part Sour Grapes, Part Political Cover, invoking “evolution” to give scientific window-dressing to a craft tradition.

But Were There Insights to be Found?

- Yes, of course. consider a few
 - new insights arrived at by *just thinking (with no funding)*
 - All layers have the same functions, a single layer that repeats.
 - Connectionless is Maximal Shared State, not Minimal
 - The dumb network ain't so dumb
 - There is only one data transfer protocol and only one application protocol
 - All (properly-designed) protocols are soft-state.
 - Networking is defined by whether Maximum Packet Lifetime is bounded.
 - There is no requirement for a global address space, or a global name space
 - Right, IPv6 has been a waste of time. . . twice
 - Multihoming and mobility are a consequence of the structure, i.e. free.
 - The resulting structure is inherently more secure.
 - Without NATs or firewalls
 - There are more and we are still finding them!
- Return on Investment?
- *FI efforts: 0/€10⁸ = 0* *RINA: n/\$0 = ∞*
 - Brings to mind the phrase, “buying a pig in a poke”

Why Did They Go Wrong?

- Treated it like a traditional engineering problem in a mature field
 - They looked at requirements. Consider an excerpt from a typical list of issues to be solved by a Future Internet Architecture:
 - Inability to diagnose problems
 - Lack of data and service identity
 - Real time processing
 - Lack of inherited and data privacy
 - Lack of data integrity, reliability, provenance, and trust
 - Lack of data integration and federation storage solutions
 - Lack of efficient transmission of content-oriented traffic
 - Lack of integration of devices with limited resources
 - Lack of effective congestion control
 - Lack of mobility support.
- These are product requirements

Source: EC FIArch Group, Mar 2011

Classic Low Uncertainty Approach to a High Uncertainty Problem.

This is *not* “just an engineering problem.”

- The issue was that the architecture had “run out of steam.”
 - Important to Remember that there is a
 - Big difference between an architecture and things built to an architecture.
 - Victorian architecture is very different from a Victorian building
 - The need is not for a new building because existing ones do not meet requirements. The need is for a new architecture.
 - Architecture is precisely that which is invariant wrt user requirements.

This Points to Incompleteness

- When an architecture runs out of steam, it points to incompleteness
 - Fundamental concepts need to be reconsidered or are not present.
 - This requires a careful investigation of assumptions, new invariants, asking the deeper questions of why, was our initial direction valid?
 - Turns out it was but we didn't take that direction (and no, it wasn't OSI).
- They were solving the wrong problem.
 - Looking in the wrong direction.
- They were asking what do we build (craft), when they should have been asking what don't we understand (science).
- But it wasn't entirely their fault
 - Failure was virtually guaranteed by the organization of the work.
- By Contrast,

The Questions We Considered

- The core architecture problems of the last 40 years were well-known and fundamental:
 - How can connection and connectionless be unified?
 - What does a complete naming and addressing architecture look like?
 - What is the nature of a layer?
 - What are the “upper layers”? Are there any?
 - Does separating mechanism and policy provide some insight?
 - What does location-dependent but route-independent mean on an arbitrary graph?
- This list produced the previous list of breakthroughs.
- And, will address the other list as well . . . and do it better

How Bad Is It?

- Pretty bad.
- In 2001, a National Research Council study told US networking researchers:

“A reviewer of an early draft of this report observed that this proposed framework – measure, develop theory, prototype new ideas – looks a lot like Research 101. . . . From the perspective of the outsiders, the insiders had not shown that they had managed to exercise the usual elements of a successful research program, so a back-to-basics message was fitting.”

Looking over the Fence at Networking,
Committee on Research Horizons in Networking,
National Research Council, 2001.

- Must be sobering to be told you don't know how to do research
- This is before all of the work on FIND and GENI.

So They Ignored the Report?

- Not at all. They tried to learn from it.
- That *is* the bad news.
 - They were trained as craftsmen. The selective pressures on them (publishing and funding) made them into craftsmen, if they weren't already. It was easier. Not taught to create scientific theory; but taught to build things.
 - When NEWARCH came up dry, they recommended looking outside networking for insights.
 - A few years later when FIND came up dry, they made the *same* recommendation to look outside the field.
 - When, of course the insights were inside the field, as they always are.
 - And as RINA has shown.

Didn't They Try?

- Of course, but . . .
 - The attempts to do theory were riven with errors, circular arguments, false distinctions, teleological arguments, etc. IOW beginner mistakes.
 - Consider their guiding “principles:” hard-state/soft-state, fate-sharing, end-to-end, etc. are all *descriptive*, not predictive.
 - This is natural history, not science.
 - A clear indicator that how to uncover fundamentals eludes them.
 - No university level textbooks; all vocational, trade school books
 - Would you ask a Carpenter to design a high-rise building?
 - We aren't training scientists, but craftsmen.
 - And if that wasn't enough, more than one of the prominent (and popular) proposals being funded aren't even network architectures.
 - and no one has noticed (or had the nerve to say so)
- Sounds pretty bleak, what do we do?

Good Question

- Engineering without Science (theory) is Craft.
- This is not uncommon in the early days of a field.
 - Early steam engine developments before Carnot
 - Faraday leading to Maxwell; Edison instead of Tesla;
 - Early aircraft attempts before the Wright Brothers and others
- Seat-of-the-Pants often gets things started but it never lasts.
 - In this case, craft was prolonged by Moore's Law.
 - “Any sufficiently advanced *craft* is indistinguishable from *science*.”
- But it is a young field, too young to have developed theory
 - Not really, early the variety of systems began the development of theory.
 - By 1975, there were at least 15 - 20 different system architectures on the ARPANET. One needed a theory to make sense of it. Theory played a key role in the development of databases and operating systems
 - But Moore's Law, the consolidation to monocultures, and the gold rush nipped it in the bud
 - And of course NIH* did its part (and its companion WDIITD).

So Now What?

- Two generations of researchers conditioned and educated in this mode.
 - We are already seeing stagnation and crisis-level problems “solved” by patch.
- The Structure of Funding Doesn’t Help
 - Narrowly Focused, Peer Review, Industry Partners.
 - How many industry partners did the ARPANET have? Zero
 - How many did CYCLADES have?
 - You’re joking, right! They Couldn’t Wait to Shut it Down
 - Would the EU fund CYCLADES today?
 - We have this brilliant idea but if we are right, it invalidates your business model. Would you give us money?
 - If want funding, then only propose what you think they will buy.
 - » A bit like how censorship doesn’t work in China.
 - Any industry partner will always be asking how would this impact my revenue. Ensures no breakthrough
 - If the work is any good, it must be disruptive.

So Now What?

- So we need to replace craftsmen with theorists?
 - No. Never many theorists in a field. There just shouldn't be none, which is pretty much what we have. A few theorists generate work for many.
- Doesn't that make it a beauty contest? Whose theory wins?
 - Not entirely, we do have Newton's Guidelines from the *Regulae Philosophandi* of 1726 (as paraphrased by Gerald Holton):
 - Nature is essentially simple; therefore, we should not introduce more hypotheses than are sufficient and necessary for the explanation of observed facts. This is a hypothesis, or rule, of simplicity and *verae causae*.
 - Hence, as far as possible, similar effects must be assigned to the same cause. This is a principle of uniformity of nature.
 - Properties common to all those bodies within reach of our experiments are assumed (even if only tentatively) as pertaining to all bodies in general. This is a reformulation of the first two hypotheses and is needed for forming universals.
 - Propositions in science obtained by wide induction are to be regarded as exactly or approximately true until phenomena or experiments show that they may be corrected or are liable to exceptions. This principle states that propositions induced on the basis of experiment should not be confuted merely by proposing contrary hypotheses.

So Now What?

- So our major research centers should become centers for theory?
 - That would be nice, but our hole is deeper than that.
 - Strange as it may seem, but almost by definition, the leading centers today, i.e. the ones that got us into this mess, yielded the best craftsmen, are the least qualified to lead us out of it, as they have demonstrated.
- OTOH, Bad theorists are plentiful. Which is why theory has such a bad rep.
 - Look at all the failed projects based on the latest theoretical fad. Those who become infatuated with an idea (often something recently successful in the market) that becomes the answer to all problems.
 - Good theorists are an odd lot. They have deep experience in both implementation and theory.
 - They tend to want to think before they build.
 - They are more apt to listen to the problem, not impose an answer.

Give It to Me Straight

- We are going to have to grow a new crop of theorists.
- We need to figure out how to select for them.
 - And reward them.
- Most current researchers tackle specific problems so will not be of much use looking at the bigger issues.
- We need to instill a new crop of students to challenge accepted views.

Academic research has lost its way.

- It is often at its best angel funding for sissies, or at worst, pointless.
 - Note the similarity between university arguments for their development offices and the incentives given to attract industry to underdeveloped areas. The benefits have proven to be mostly one way (see Jane Jacobs). Why is academia any different.
- Academic research should be targeted at the long term, not cheap (and subsidized) product research, working out the implications of theory, improve our understanding, simplify what needs to be learned, create challenges to old theory.
 - Students should come out ahead of industry not behind it.
- Good theory Improves and Simplifies what we teach.
 - Imagine teaching Electrical Engineering without Maxwell!
 - Diversification means more courses, which means more revenue

If We Don't Change Our Ways

- I am Worried. Which is Why I am Talking about this.
- Our Research is Stagnating
- Thanks to Globalization, We Will Go the Way of other Civilizations.
 - This isn't an issue of West or East or North or South.
- It is *us!* All of us!
- And Too Much Depends on It
- We Must Break the Trend!

Okay, Enough Generalities

What Should We Be Doing?

- Well, I Know it When I See It! ;-)
- I Can Tell You What I Do, I am getting better at understanding it.
- For a long time, I would talk about doing the algebra before the arithmetic, or good designers had the topologist's vision defect.
- Maximizing the invariances and minimizing the discontinuities.
 - Abstraction is finding the right invariances
 - If the devil is in the details, you did something wrong
 - The wrong invariances, there was something that wasn't understood.
- Still pretty general, Day, get specific!
 - Okay, but one more, no, two more things:

Just Two, I Promise

- As I have said, Learn to Listen to the Problem
 - Sounds crazy, but the problem will tell you what to do.
 - Find the answer, the problem dictates, not the one you want.
 - Kettering said, “A Problem Well Defined is a Problem Half Solved”
 - Collecting models (reading widely) helps to guard against . . .
 - When all you have is a hammer, everything looks like your thumb! - O’Dell
 - Always look at a problem from the point of view of the organism, not the observer. – von Forester
 - The Problem with Reductio Ad Absurdum is Knowing when to Stop!
 - See “Rosencrantz and Guildenstern are Dead” - Tom Stoppard
- And a sense of aesthetic really helps. Well, okay, 3 and a half.

Examples

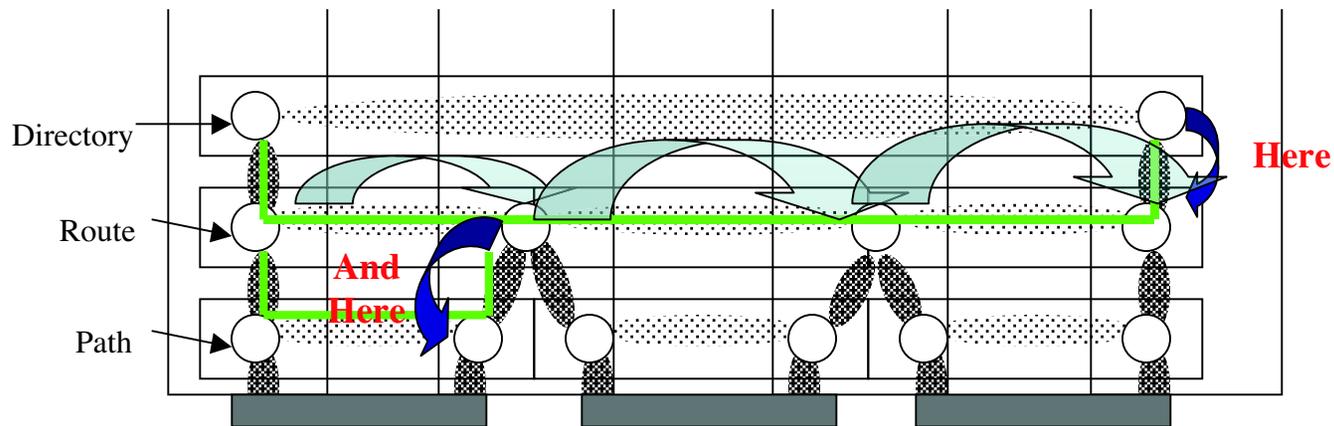
- How Did I Come Up with What is Now called RINA?
 - Not by setting out to do new architecture or build anything.
 - Although I always have an implementation strategy in mind. ;-)
 - But trying to figure out what I knew (and didn't)
 - Try to understand “all that was the case” without concern for politics or what was already there. (Once know where to go, easier to figure out how to get there.)
- Separating Mechanism and Policy in Protocols
 - Separates the Invariants from the variant.
 - What did We Learn?
 - Loosely coupled and tightly coupled mechanisms
 - TCP had been split in the wrong direction
 - Fits hand in glove with Watson's results for a much cleaner implementation
 - That is much faster, lower overhead.
 - It took a long time to admit they were really distinct protocols.
 - Even longer to admit there was no distinct protocol like IP.
 - But the problem kept insisting. Do what it says!!

Watson's Delta-t

- In 1978, Watson proves the *necessary and sufficient* conditions for synchronization for reliable data transfer is to bound 3 timers.
 - SYNs and FINs are unnecessary
 - Whoa! Where Does that Come From?
 - Early work tried to avoid time-outs. Found out they were unavoidable.
 - Okay, then turn it around, If need time outs and what else?
- Then notes, all connections exist all the time. The TCB's are merely caches of state for connections on which there has been data recently.
 - A real particle physics view Cute!
- What this is saying is, decouple port allocation and synchronization.
 - *This* minor point turns out to have major implications:
 - greatly improve robustness and improves security.
- It also defines what is Networking:
 - If MPL can be bounded, it is networking; otherwise remote storage.

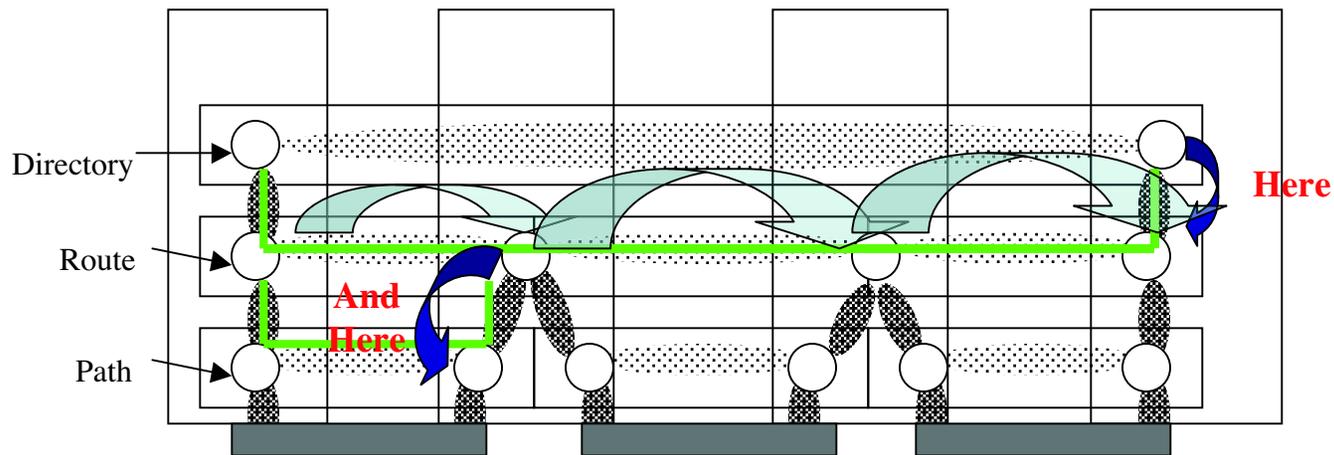
Carefully Reading Saltzer for the Umpteenth Time

- But also buried in here is a fundamental misunderstanding: Layers are different in Networks than in Operating Systems.
- The dominant property is Distributed Shared State of a Given Scope.
- Then routing should be on the address where the relaying is?
 - Not on the address of the Layer below. Well, duh!
- A Problem we are still living with: The Fundamental Flaw in IP.



Carefully Reading Saltzer for the Umpteenth Time

- But also buried in here is a fundamental misunderstanding: Layers are different in Networks than in Operating Systems.
- The dominant property is Distributed Shared State of a Given Scope.
- Then routing should be on the address where the relaying is?
 - Not on the address of the Layer below. Well, duh!
- A Problem we are still living with: The Fundamental Flaw in IP.



I Could Go On and On

- My Colleagues can tell you about the calls that start with “You are not going to believe what fell out . . . “
 - No new concepts have been added to the model for at least 7 years, but many new insights keep falling out.
- The fact that a global address space is unnecessary was a consequence of the model that was there for two years before we saw it!
- The IPC Model itself was created as a teaching tool, but was the unifying principle that explained the patterns we were seeing.
- Security was not considered at all, yet the model naturally forms a securable container and does it with less complexity.

- And this stuff is a heck of a lot of fun!!

So Is that the Answer?

- Elepharhino!* I have been trying to practice what I preach!
- What I have touched on here is where *patterns* derived from fundamentals have taken us. It does have promising properties.
- But like any theory, it should be explored and tested.
- Is there a simpler model with more predictive breadth?
- Who knows? The fun is in the hunt!

*The answer to What do you get when you cross an elephant and a rhinoceros?

Very old (and bad) joke

So Come On!

- Lets go back to being Scientists!
- Find neat things no one knows!
- If the old guard doesn't think you are crazy, you aren't working hard enough.
- When I was Your Age, Anyone my Age Didn't Know Squat!
- And We Were Right! So What is Wrong With You Guys!
 - Good grief! Where is Oedipus when you need him!?

Oh! & BTW

- If we had built on CYCLADES,
- It *would* be “just an engineering problem,” but it isn’t.
- How So?
- Left as an exercise.
 - Hint: If the Internet were an Operating System, it would be DOS.
- Now it is a whole lot more than engineering.
- The world depends on us to get it right the first time without stumbling.
 - Not exactly something you want to leave to an artisan tradition