Digital Image Design Incorporated

Perspicands

Research, Applied: Working Innovations, Techniques and Products

Per·spi·ca·cious: having a ready insight into and understanding of subtle or complex things

Per·spic·u·ous: (of a representation) clearly expressed, easily understood; lucid

Per·spi·cand: (our back-formation) the easily-understood representation

Contents

1 Research into Cognitive/Analytic Tools: Perspicands *2*
2 Applied to Real-World Tasks *26*
3 Background: Key Participants *50*

1 Research into Cognitive/Analytic Tools

An annotated portfolio of techniques that let visual tools extend the human mind

A New Kind of Tool

Illustrates the task itself, not the rules of the tool

Knowledge work is just starting to take a significant step forward thanks to a new kind of tool. Risk can be reduced; speed and quality of insights and actions can be improved; creativity, morale, and satisfaction can be raised; training and support costs can almost be eliminated.

These benefits result from combining hundreds of small techniques that arise from a single fundamental shift in perspective: make the tool *illustrate* the task—fit it to the cognitive steps the human mind does to accomplish the task. We do this rather than forcing people to learn all the idiosyncratic ways designers have coerced the task into form-like widgets, then try to map their learnings back to there real thing they're trying to do. This indirection hobbles thought.

Every task we've studied has a grounding in space. Even (perhaps especially) the most abstract thought is grounded in spatial thinking. Visuals don't just *represent* information—when they illustrate it well they *become* it; they contain everything needed to do the task—or even teach it.

My collaborators and I have quietly deployed dozens of these tools over the last two decades. You're likely to have seen one or two at work at the New York Stock Exchange or other major financial institutions, featured in the New York Times, Nature, Geo, ID Magazine, SlashDot; or on the wall at MoMA or the Whitney Museum of American Art. Proven to work, it's now time for us to broaden their use; it's time to help more people think more effectively.

We gather many techniques and representations here, in categories identifying the kinds of thought processes they support. Innovated, honed, tested, deployed, and working; we're now stretching this new approach to toolmaking; addressing today's most critical human/data challenges. Today's common labeled-rectangle approach to interface design introduces a false complexity. Once you can *see* those challenges they're often not as complex as they seemed.

Brad Paley · brad@didi.co

"I bring evidence to show that spatial thinking, acting in the world with the things in the world, is the foundation of thought; not the entire edifice, but the foundation."

—Barbara Tversky, Mind in Motion

"If I pursue a beam of light with the velocity c (velocity of light in a vacuum), I should observe such a beam of light as an electromagnetic field at rest though spatially oscillating."

—Albert Einstein's spatial thought experiment, Autobiographical Notes

"We knew that DNA was important… But we didn't realise I think just how important it would be. Put in other words we didn't realise that the shape would give us a clue to the replication mechanism."

—Francis Crick, BBC TV interview

Dashboards

Overviews let us get our bearings and prioritize—but they must be carefully tuned, not just pulled together: they give us the lay of the land best when fitted to specific tasks and cognitive sequences.

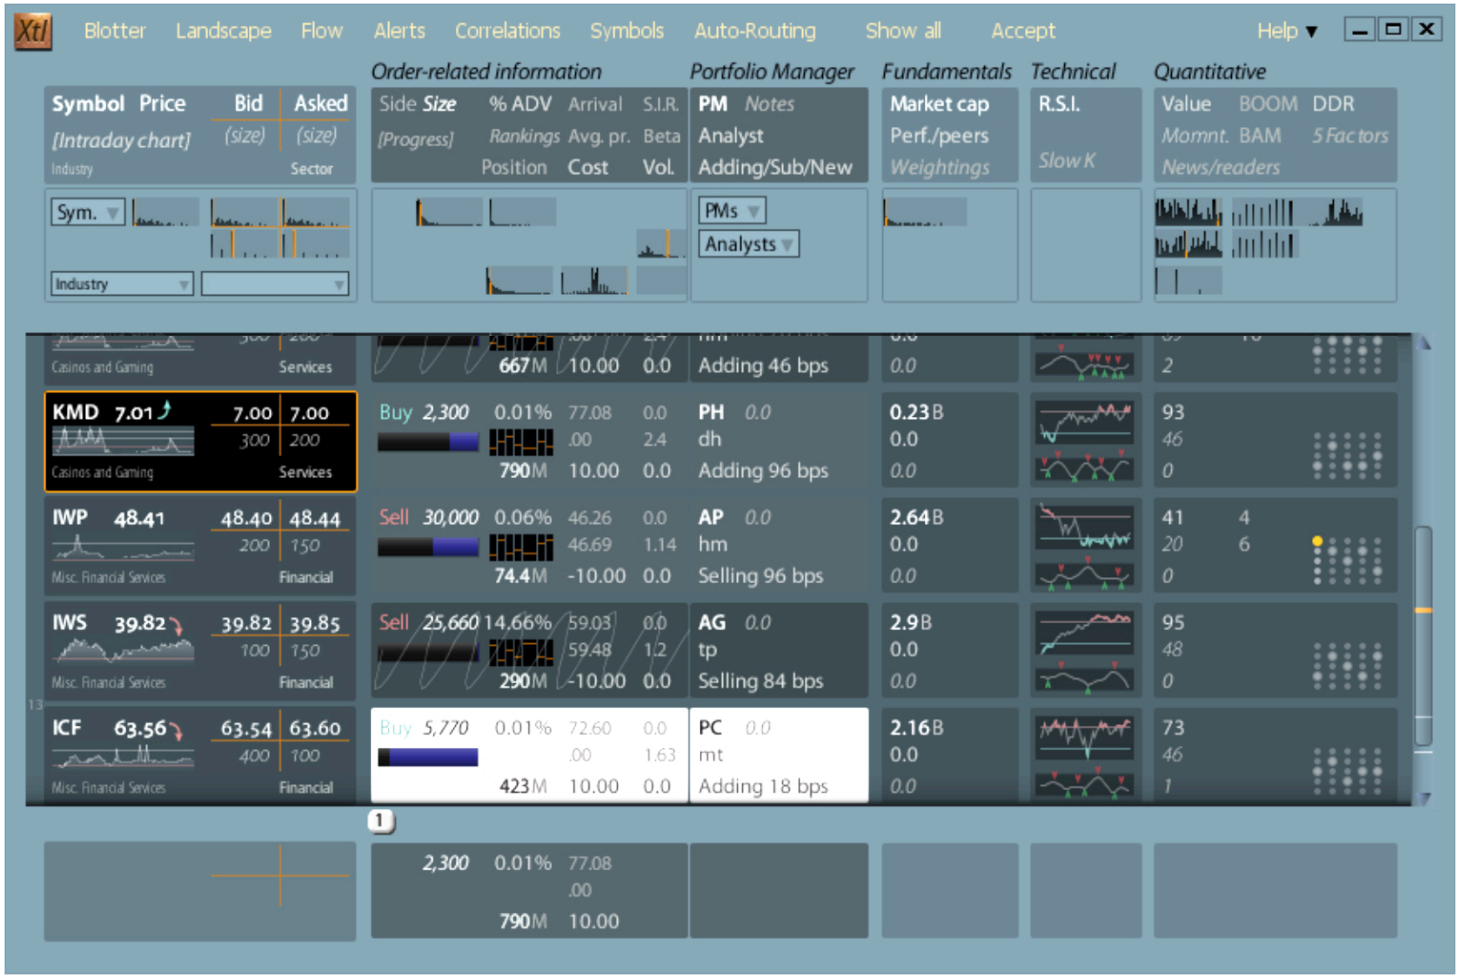

This equities trading "blotter" view groups information by intellectual context. 50+ columns worth of data, needing four times the width shown here when displayed as numbers, were arranged so no horizontal scrolling was necessary. This was absolutely critical to traders who needed to get an order working in the market within 30 seconds. Before us the off-screen columns were hidden and it was difficult to get to them, then return the view to its usual, working, position. Worse, in a situation where a trader might be asked why they didn't take that information into account, they were open to a "but it was right there in front of you" attack—when in fact functionally it was completely unavailable.

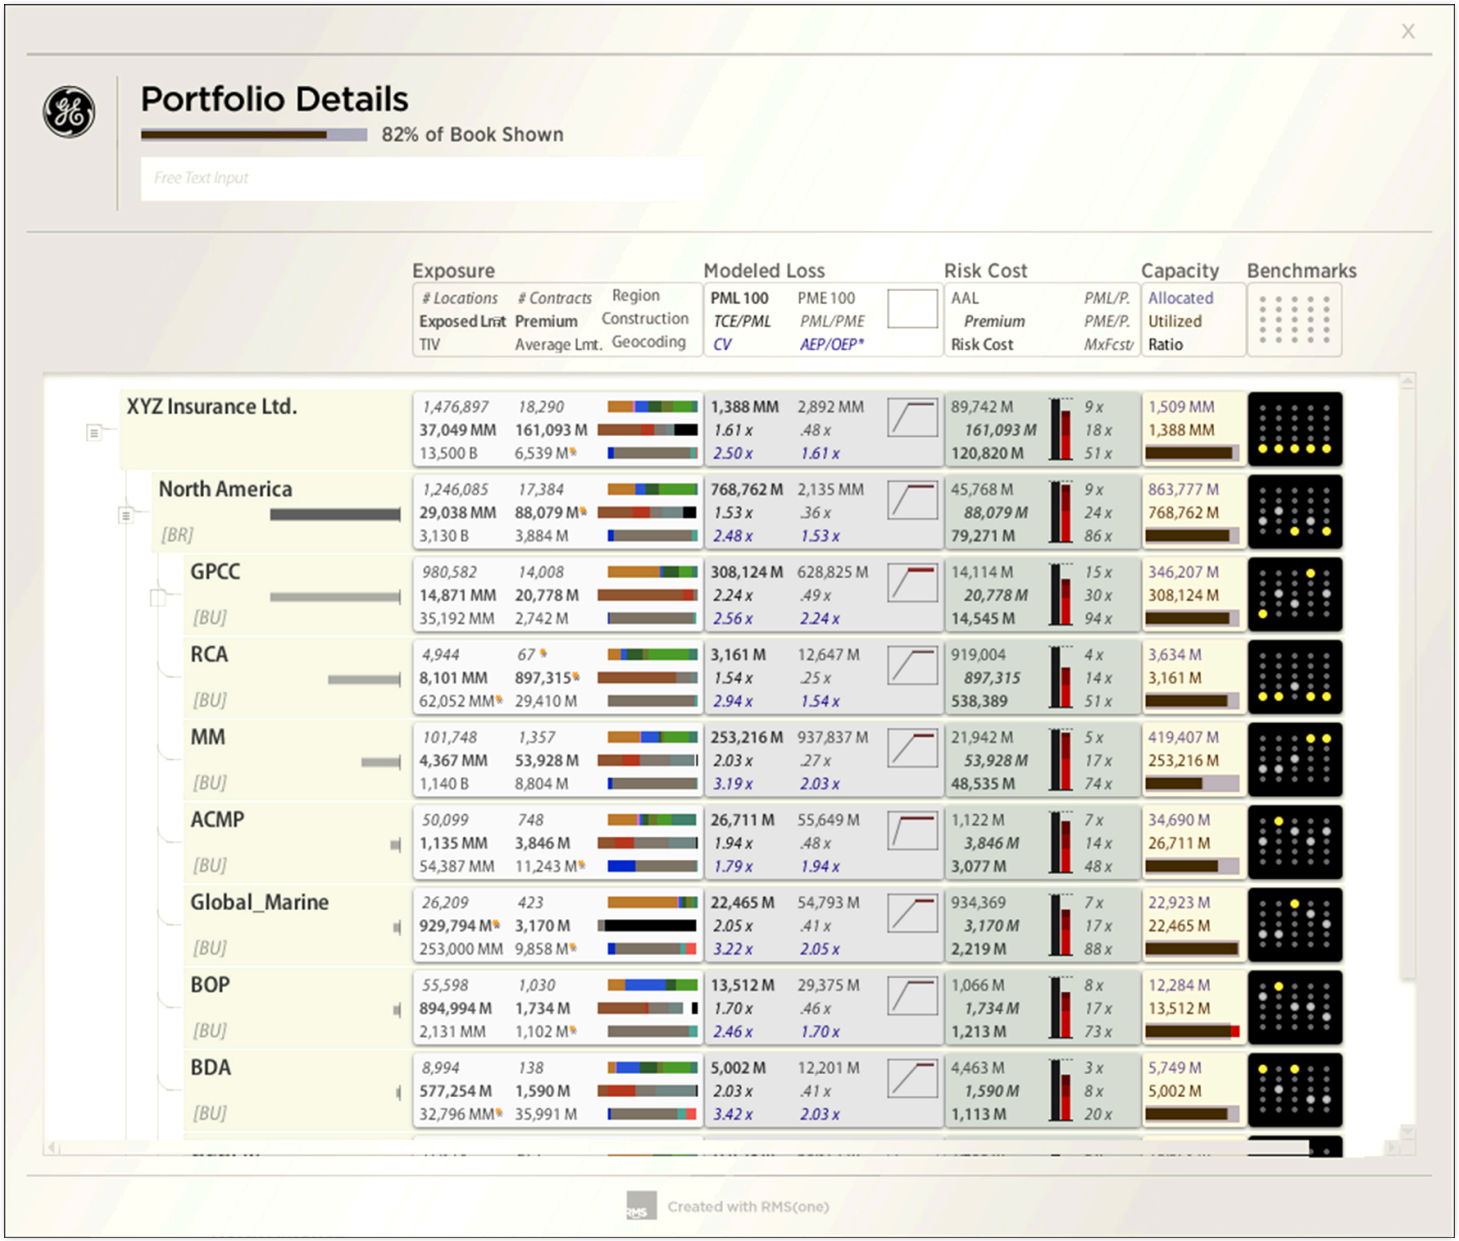

This risk-management portfolio viewer from the field of re/insurance also groups information by intellectual context.

In general, spreadsheet-like grids and blotters are incrementally expanded to become important decision-support tools; they end up with all the data one needs to choose the next task—but they're rarely sufficient to complete it in an optimal manner.

Compressing numeric data with visuals in the cells, and making it more easily absorbed, saves space, time, and most important: cognitive effort. The "real" numbers behind the visuals are often critical; when they are we put them just a click or hover away.

Transaction Support

After the overview: focus on a single task. We're more effective, more comfortable, and get into flow state more easily when a tool fits each of the sub-second cognitive steps in a task.

Forcing experts to do every task with a scattered set of apps means they have to find, then integrate all the data they need—*every day, every time they do the task*. That should be the *tool's* job. Our tools integrate it all exactly *once*, using microservices and interop tech. More important: we fit tools to a fine-grained cognitive deconstruction of what's being done. Since this is done before any code is written, the effort is properly factored out, done just that once during a project's non-recurring engineering phase. The effort of integration never has to be done in production, e.g., at the trading desk.

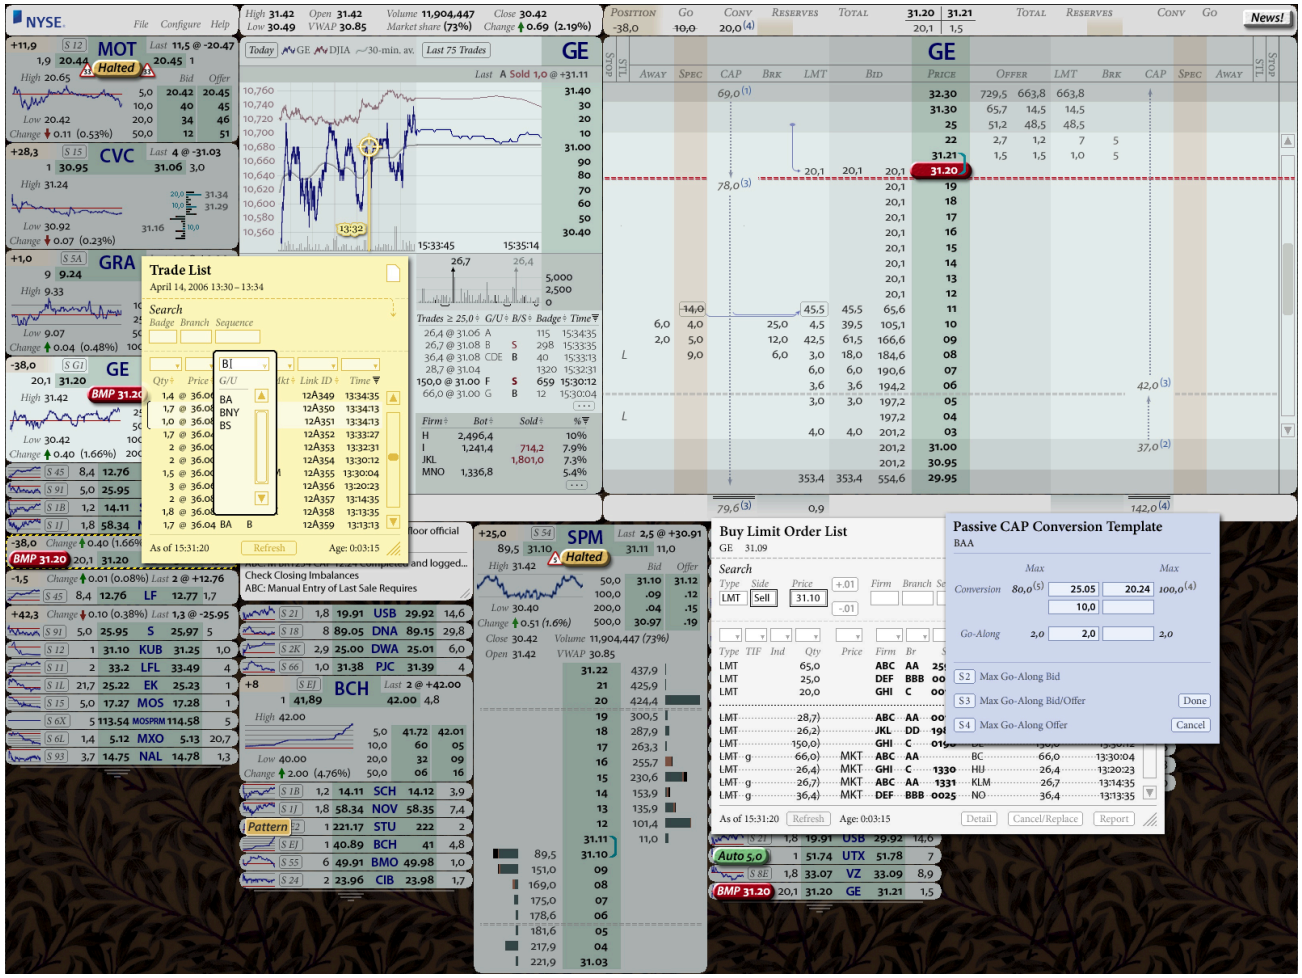

A NYSE Specialist's workstation shows 4x as many stocks as before, so 4x the characters/digits are on the screen—and it shows data for some 15x as many *actionable insights*. That's a better metric for display effectiveness; harder to quantify, but it addresses what the interface is actually *doing* better than, e.g., hardware-centered mouse clicks or design-centered information density.

The larger multiple for actionable insights comes form the fact that each tiny time-series line we added allows experts to make dozens of inferences. We also enabled easy inter-time-series comparisons by standardizing on drawing one horizontal tick line per *percentage*—not dollar—-change. An intentional side effect is that the heavier/busier charts, charting stocks that moved through more one-percent ticks during the day, are the stocks that people typically care about more: the visuals draw the eye to the business priority. The red line is the opening price: another business landmark/priority.

Our Active Order Landscape ranks orders by expected cost (vertical) and market footprint (horizontal), so traders can easily find the critical trades and work them: they're all in the upper-right quadrant. Better than any ordered list: this organization lets them make intelligent trade-offs. If more decision parameters are needed we can put more on the tiny order tiles: they're not just marks; here they show if there are red/amber alerts (see next page) and execution progress. Upon seeing the prototype one trader said "I don't need anything else you designed—this lets me do everything."

Double-clicking an order above brings up our Order Handler tool; it ranks *all* possible venues by *all* relevant criteria, and lets you instantly see how one venue (here, COWN) plays out with respect to every criterion, even if it's outside the scroll window (by the orange tick in the scroll bar). Unheard of before we invented it.

Our NYSE handheld was independently tested to make "refills," a broker's most common task, 20x faster: 7–10 seconds to *½ second*. Tap an order (yellow, *top left*) and an overlay window shows child orders with curves to where they'll execute. Then drag a template (*middle:* dotted outline + blueprint colors means "not real yet") into the market (*bottom*). Drop position determines the price, so one can "feel" this order in the context of the market before committing.

Alerting and Association, In Context

Problems aren't one-dimensional–and they can live anywhere.
A common strategy tries to order all alerts in one list.
The natural world does it differently—and far better.

In the *real-world* jungle (one primary environment our sensory/perceptual systems evolved to handle) anything unusual might kill us; these "alerts" can be anywhere—*not* in a list—but we react instantly.

Left: some kinds of alerting needs many lists: one for each critical decision dimension.

Below: Traders can set and drag amber- and red-alert thresholds as the market changes in real-time. Dotted guides show standard deviations and their backed-out normal distribution curve (the usual way alerts are set)—and how poorly that matches the real-world data: almost never distributed in a normal curve in trading. This example sets yellow alerts just inside a market-related cognitive barrier at ±$1.00 (note many trades clustered there), red alerts just outside it.

A set of word tags points into the tamer, everyday jungle of the expert. (Here, a portfolio manager: all of their portfolio's positions are sorted into hierarchical buckets, colored by market sector.)

Thin blue tendrils show every position to which the tags are applied, giving an instant feel for their distribution in this known "ground truth" view, one of the PM's home screens.

When hovered, a tag draws all the way to its tagged referents. If the PM wonders which tags apply to a sector they're considering now, they can scan the cursor over the entire list of tags in a second or less, coming back to the places they saw the most lines.

We engage a mental apparatus called "covert attention" in the visual perception literature: the PM is attending the cursor, but primed to react when tendrils touch the sector they care about.

Even the stubs of tendrils from un-hovered tags give a first approximation of how applicable each is to the portfolio and roughly where.

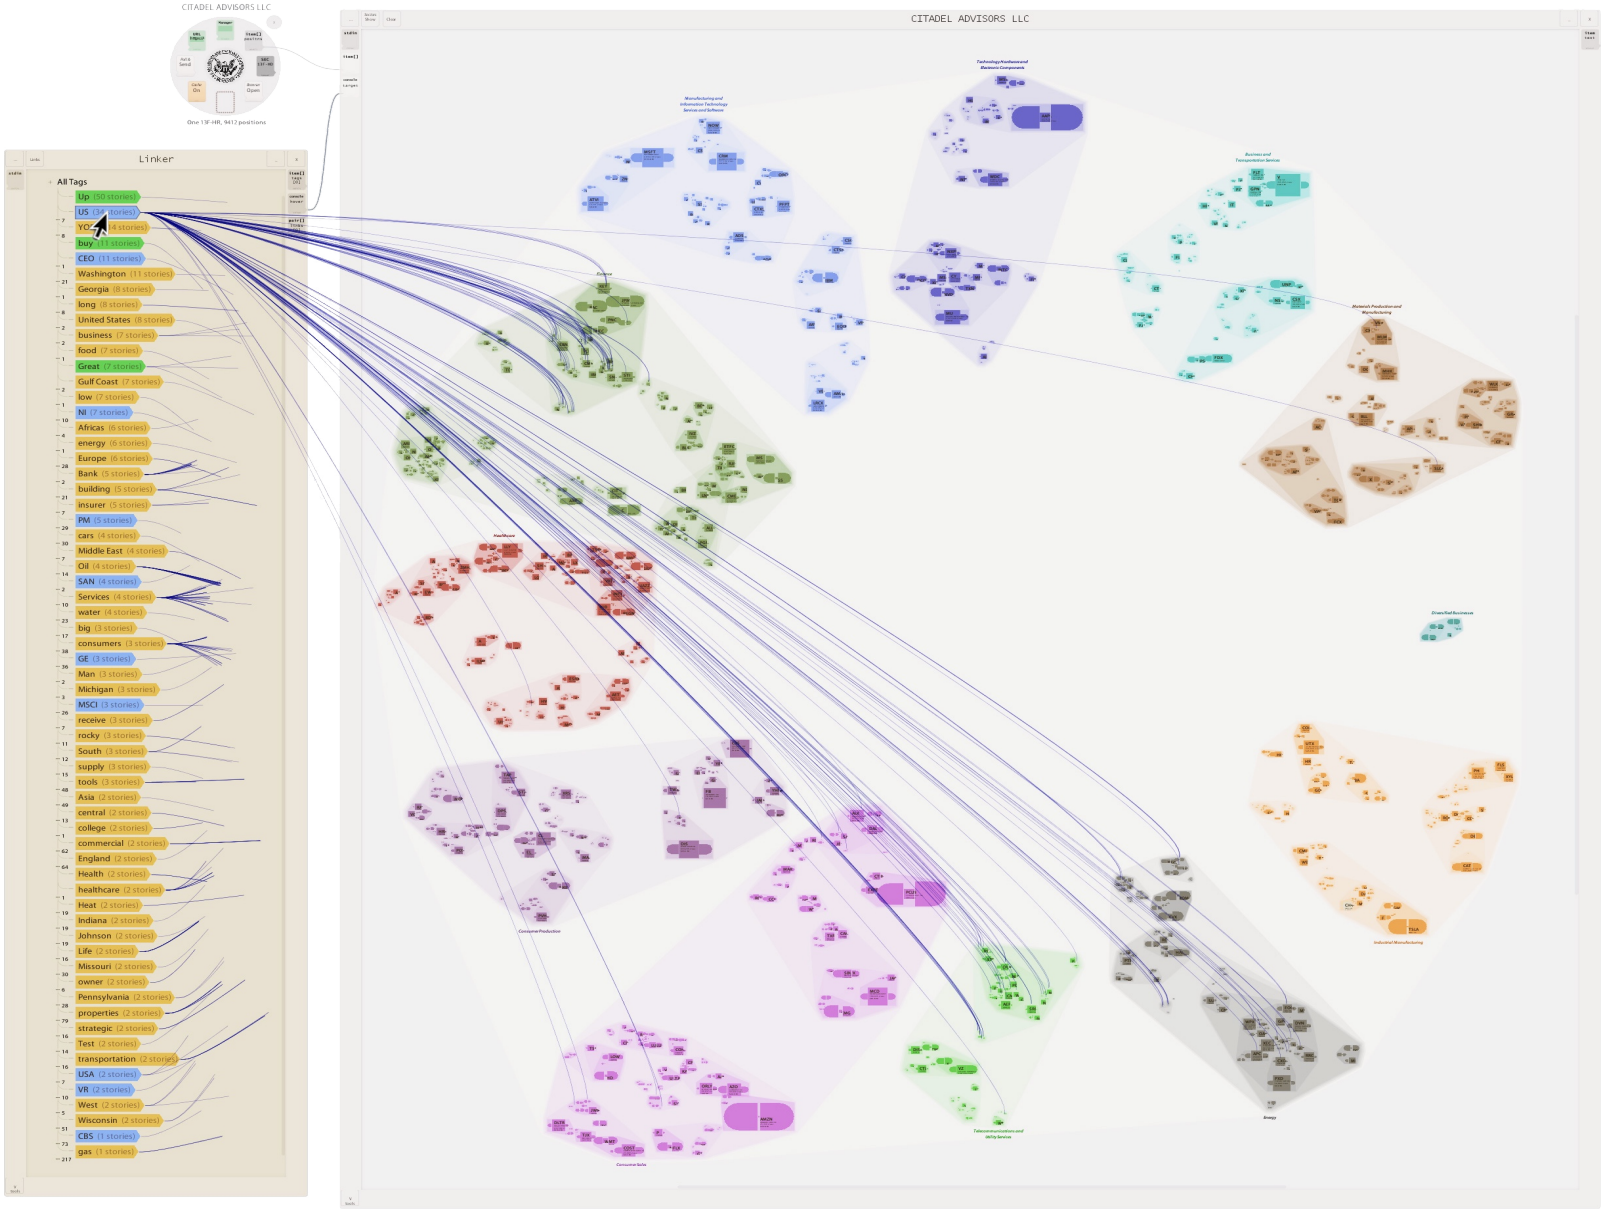

Visual Analytics

The eye supports the mind; lets it think more deeply. Manipulating data, e.g., in multiple dimensions—and seeing it that way—lets people find higher-level patterns.

People can only rehearse 7 ± 2 chunks of information in short-term memory; we can only synthesize four in working memory. And there's a huge cost to swapping in a new chunk: if you have to read or navigate to find it you may lose two or three. But when *visuals* act as a kind of second-level mental cache this swap is instant. Cognitive load is reduced, leaving more of your mind free to generate insights—or think with collaborators or clients.

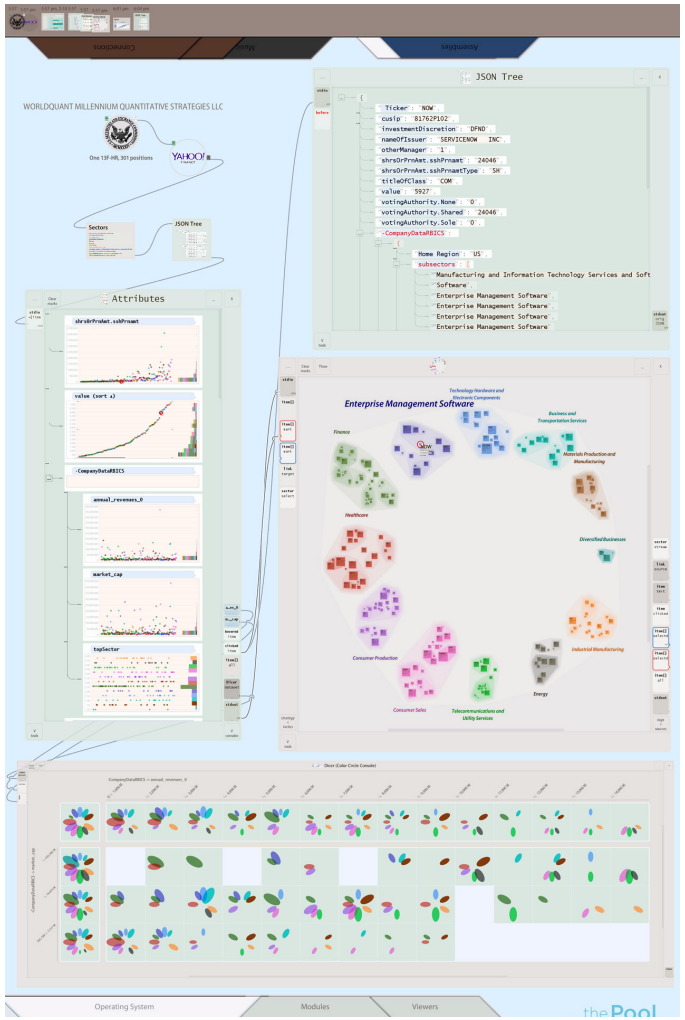

Our integration+analytics+collaboration+presentation+deployment environment, *the Pool*, lets data and ideas flow: from input to transforms to filters to visuals and back out again. Every step in a complex process be worked out and repeated, watched as it operates on streamed data, understood by others. A visual script.

Left: A 13F-HR holdings report flows through our AttributeView to feed a 2-D table in our Dicer, an *n*-D data-cube viewer. *Right top:* a 4-D table, with a bar chart in every cell—each cell can be an entire task-specific mini-report. *Bottom:* a 5-D breakdown of an insurance portfolio.

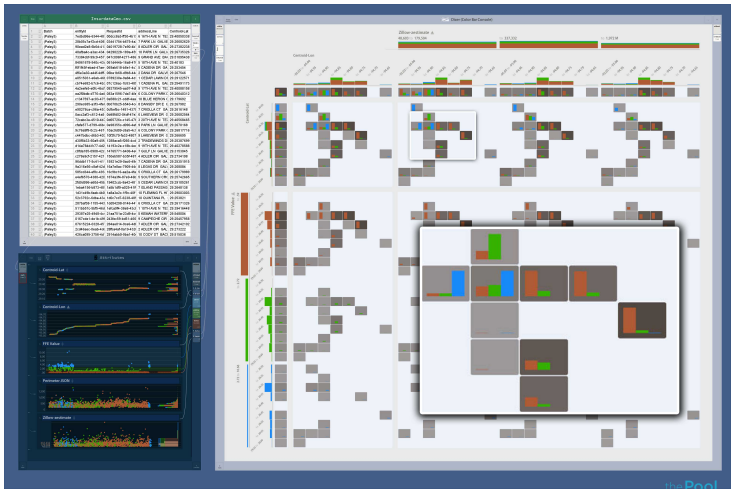

A bit of history. We had a team of five doing some of the first innovative, rigorous financial data visualization at Lehman Brothers, way back in 1989. Thought the CTO hired us to apply the hot new tech of VR, we also fell back to a tried-and-true tool: easily shared & annotated paper.

Our first draft (top): try to find clusters of + marks among the circles—they indicate where an early neural net's trades succeeded. (The pre-before they started with was worse: one had to find + marks among - marks!) *The final tool:* trade results (this of them as impulses: spikes normal to the plane of the paper) are convolved with a gaussian filter (because proximity in this strength-of-recommendation vs. trade-result space mattered) and summed, then displayed using white for success (positive results) and black for failure. The convolution additionally captures the *magnitude* of success or failure, lost in their initial work.

This image even taught us to trade *against* the net's recommendation where it's blackest: the net is so *consistently* wrong we can bet against it! Something we weren't looking for.

Statistical queries only return what you ask for—visuals let you keep thinking.

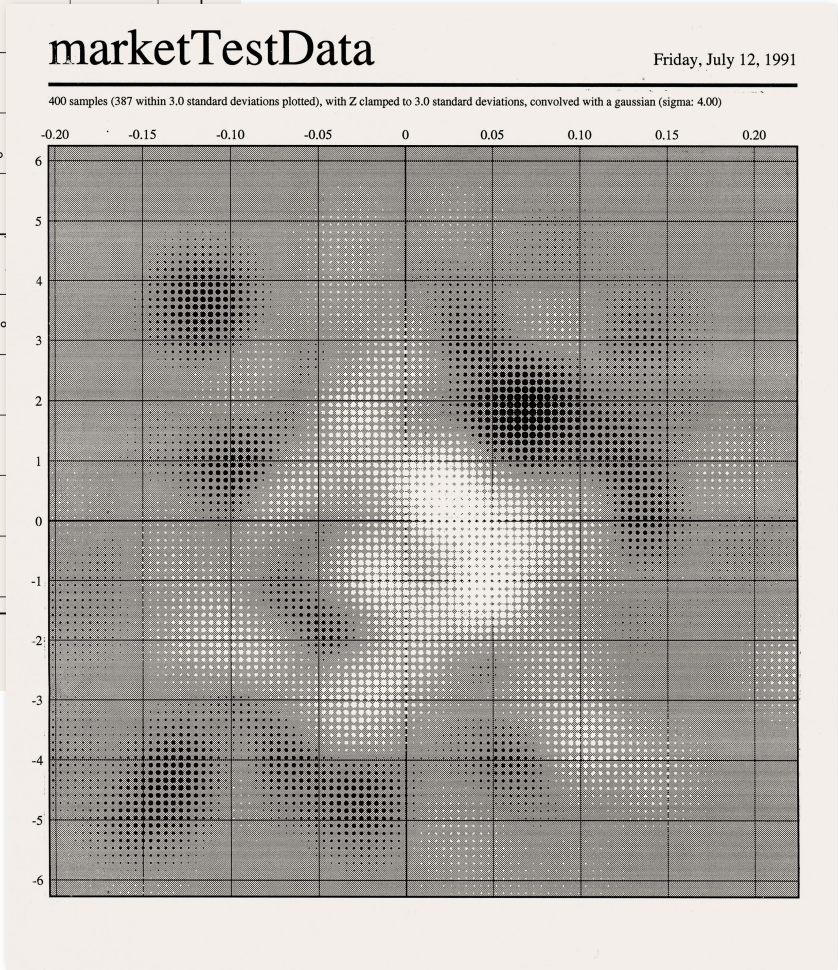

Perceptual-engagement note: we use small circles on gray rather than a continuous-tone heat map to create a reading that's both qualitative *and* quantitative. People can match the size of the circle to a size key (not in this version) more accurately than judging a color that's next to other colors. In fact all heat maps we've seen have this serious problem: the eye reads color *with respect to its immediate surroundings*, so you can't accurately compare similar colors if they're separated by other colors. This means you can only be assured of finding a *local* maximum in close cases. Here it's easy to find a global max or min: the circles are largest there. It's even easy to compare global max and min: the biggest white dots are slightly farther from zero: larger than the biggest black ones (less gray shows through).

Hypothesis Formation

Brainstorm with visuals—but don't commit. The eye is fooled more easily by visuals than statistics. It's powerful—but it's the embodiment of 510,000,000 years of evolutionary kludges.

This is *not* to say we shouldn't use visuals! They provide an unparalleled scaffolding for higher-level thought. But those thoughts need to be researched and tested using less idiosyncratic methods. Many of the eye's biases can corrected for (e.g., by TextArc's multi-attribute layering, or the bright gaps on the upcoming What-If page, 19); doing that requires a deep understanding of both eye & mind, not typically taught in a graphic artist or coder curriculum.

Critically important: if we want to be rigorous—to support anything like the scientific method—visuals must be tightly integrated with *verification* tools.

To perform properly in a workflow they also need to communicate with with tools that let people take *action* on what they see. Stand-alone visualizations are the norm. By definition they neither confirm nor commit; a dead end, not part of a domains real work process. We need to integrate them so we can both *trust* and *use* what we see.

TextArc displays a text radially and clockwise around an arc. It also shows words and distribution glyphs, here: one pair for each of the 2,500+ individual word stems in *Alice's Adventures in Wonderland*. Their star-like rays point to every place the word appears in the text.

If a Structuralist analyst hypothesizes something like "Are *Hatter*, *Dormouse* and *March Hare* all distributed like the word *tea* because they're at a tea party?" they can verify/falsify with a test. (For textual analysis, that test can be as quick as reading the text, right inside the tool.)

Another distribution comparison: why do *King* and *Queen* appear in almost all the same places—except four thin lines emanating lower right from *Queen*? Reading the text show us that's what foreshadowing looks like. She doesn't actually appear yet, she's just mentioned in advance for a more dramatic entrance.

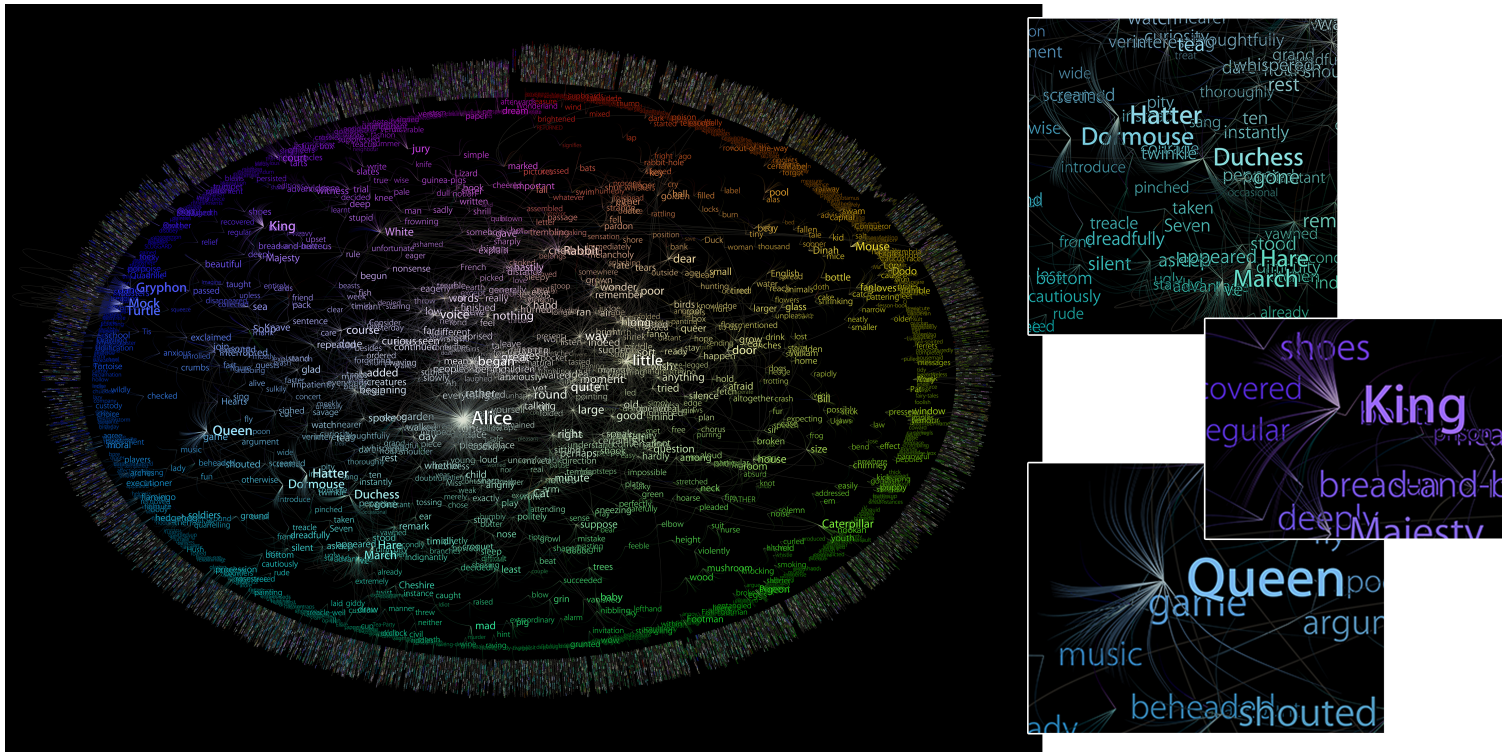

TextArc's layout, glyph, and layering can be generalized and applied to any dataset, not just text, where the goal is understanding thousands of individual elements (think of tags or attributes) with respect to a single linear ranking.

Causality

What will follow—as the practice of analytics evolves. Causal models: stronger than correlation, but we're just starting to understand them. Intervening helps.

Judea Pearl developed Bayesian networks, but when he found he couldn't get them to incorporate the concept of cause and effect he left them to his graduate students. They only quantify correlation, they changed the world.

Pearl went on to develop Structural Causal Models—and they're starting to change analytics again. Why? Killing the rooster will in fact *not* stop the sun from rising—but most statistical models can't tell you that. They show only correlation, leaving analysts to *guess* what actually drives the relationship. SCMs specify and quantify *causality*.

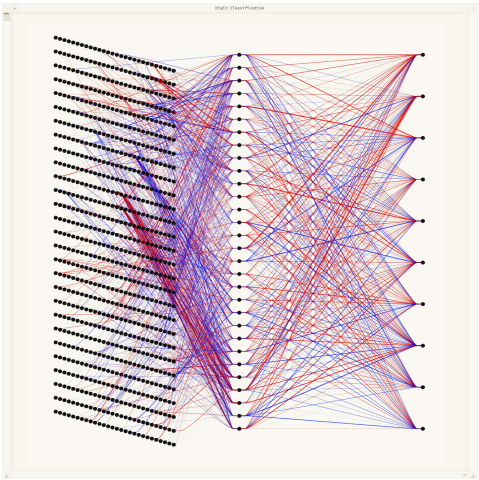

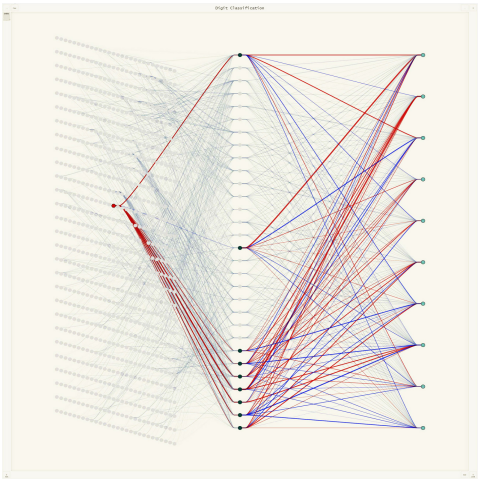

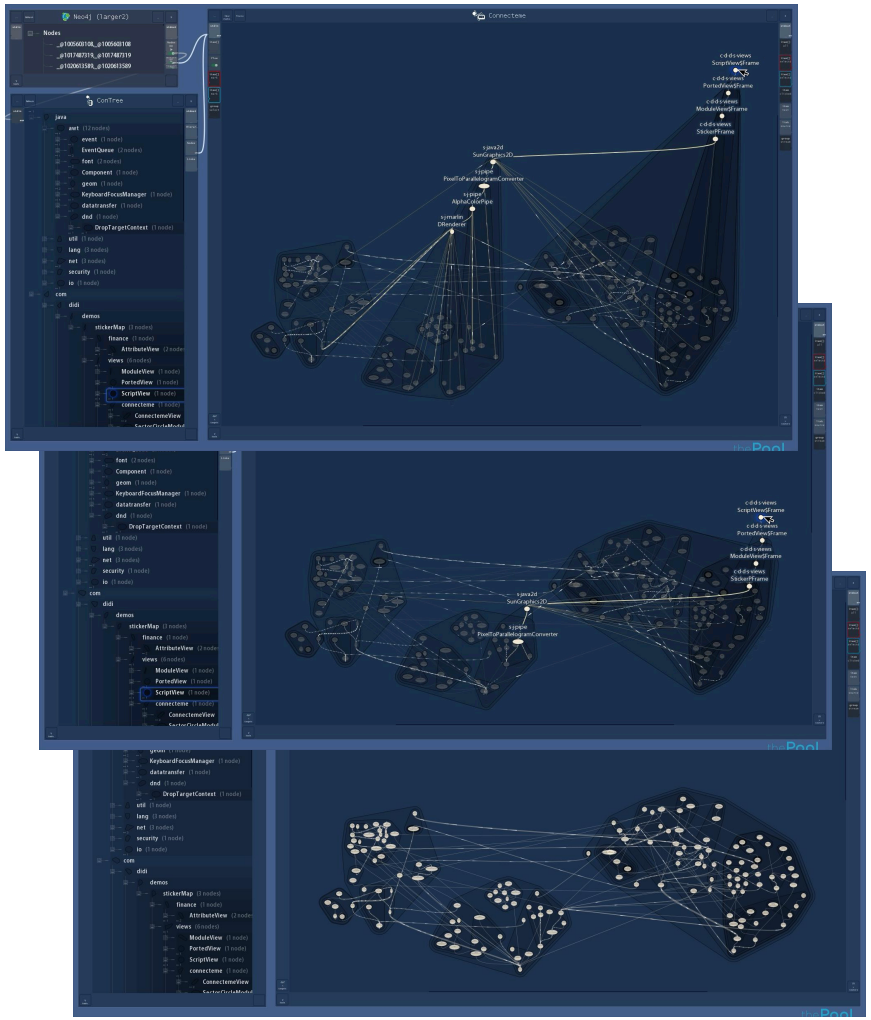

Above: Pulling "a strand of causal pearls" out of a network (bottom to top)

Network *Understanding*

GraphVis can help, if humans are allowed to direct the layouts.
Graphs can encode every part of a complex issue as nodes/edges.
But meaning is hard: semantics is human—we need to add it.

People learn better when they can play with and organize things themselves, as compared to passive viewing. This is well known in pedagogy and cognitive science. Our tools let people help arrange data to add in the meaning they know. (E.g., put the C-suite on top, infrastructure bottom.) This supports sense-making and evokes useful/actionable insights. It's hard to absorb/use data based on robot-arranged topology alone.

Algorithmic layout rules are an attractive fallback when semantics aren't available, and create lovely organic-looking slides, but they often foreground the *rules of the algorithm* more than the data. If domain experts can put in the data's meaning, they can join with the layout algorithms as intellectual cyborgs when their top-down organizing skills are supported by bottom-up algos: people move the big stuff, robots tidy up. As Cognitive Engineers we design tools that work at the level of meaning by finding domain experts, then fitting the tools to their best practices—the tools are informed by meaningful real-world processes and goals.

Next Page: global supply chains: ~6,000 nodes & ~30,000 links, but still readable—and NetMap still evokes insights and allows drilling in for details. (Example: who supplies Apple and how they stand with respect to sector peers)

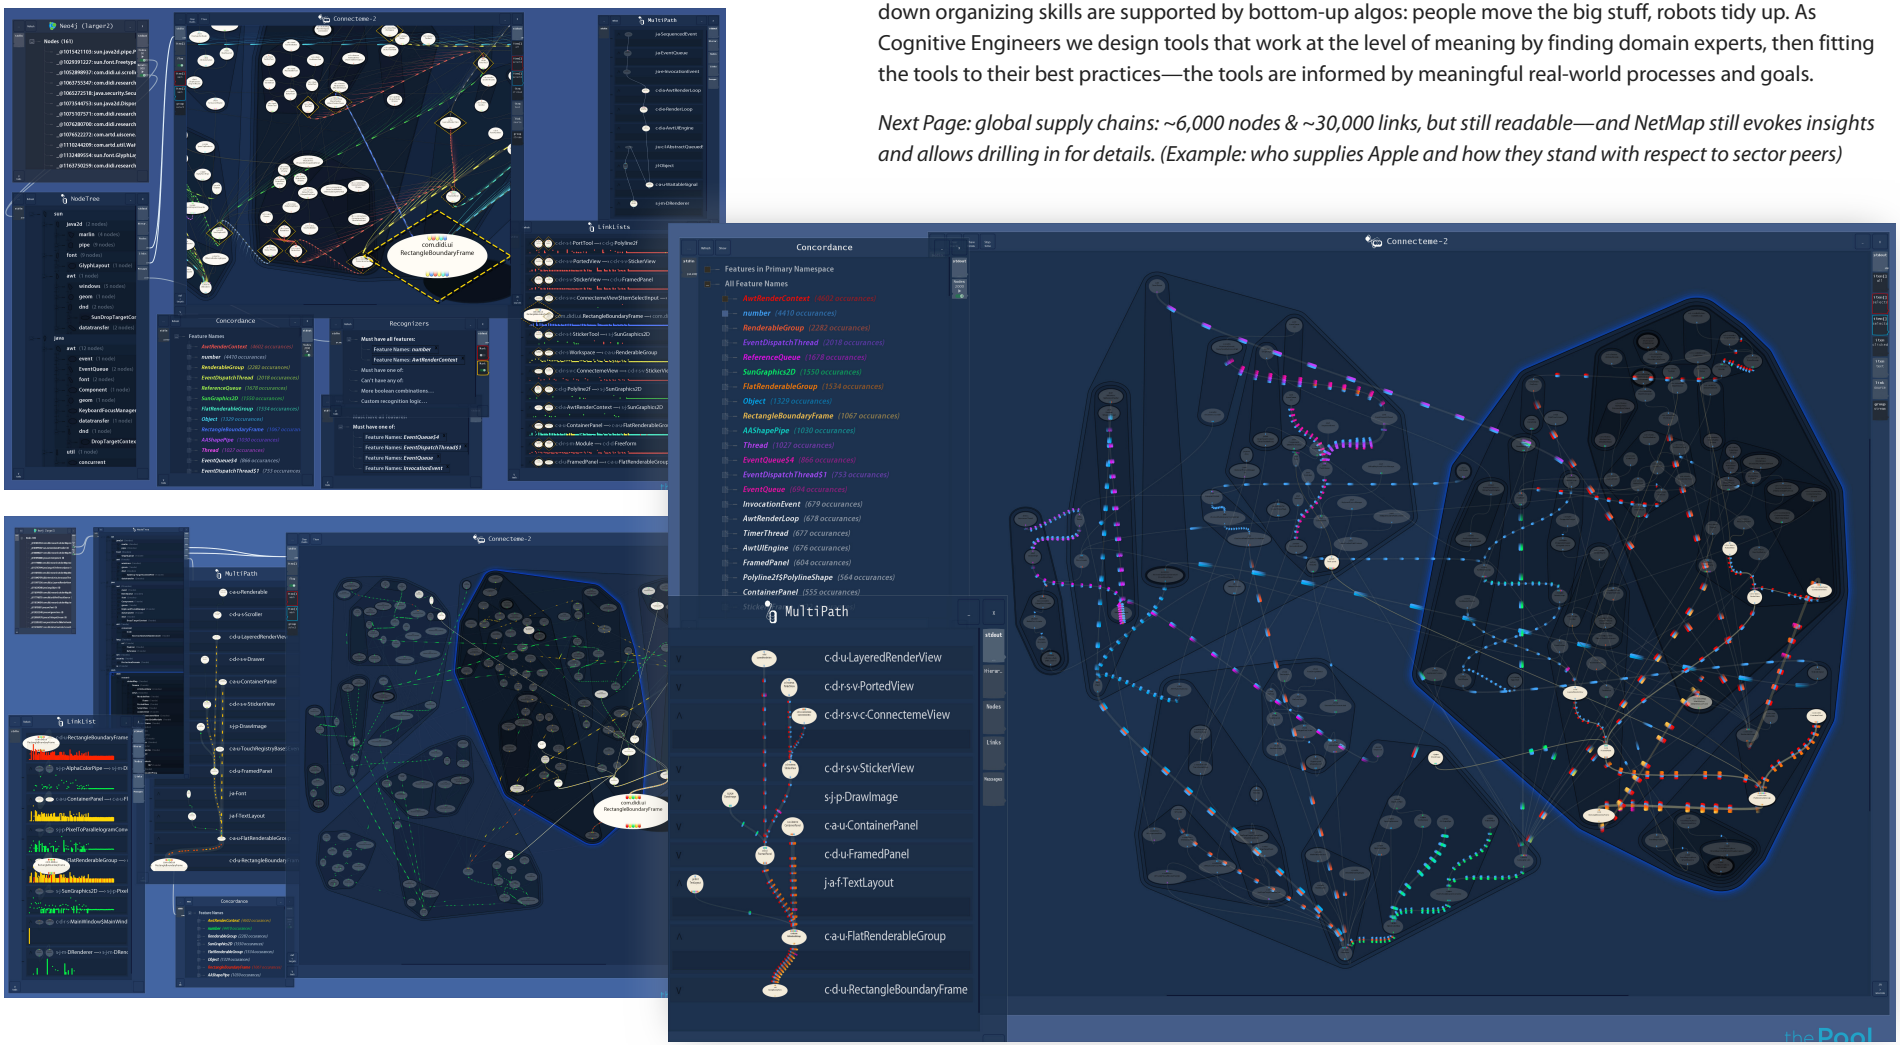

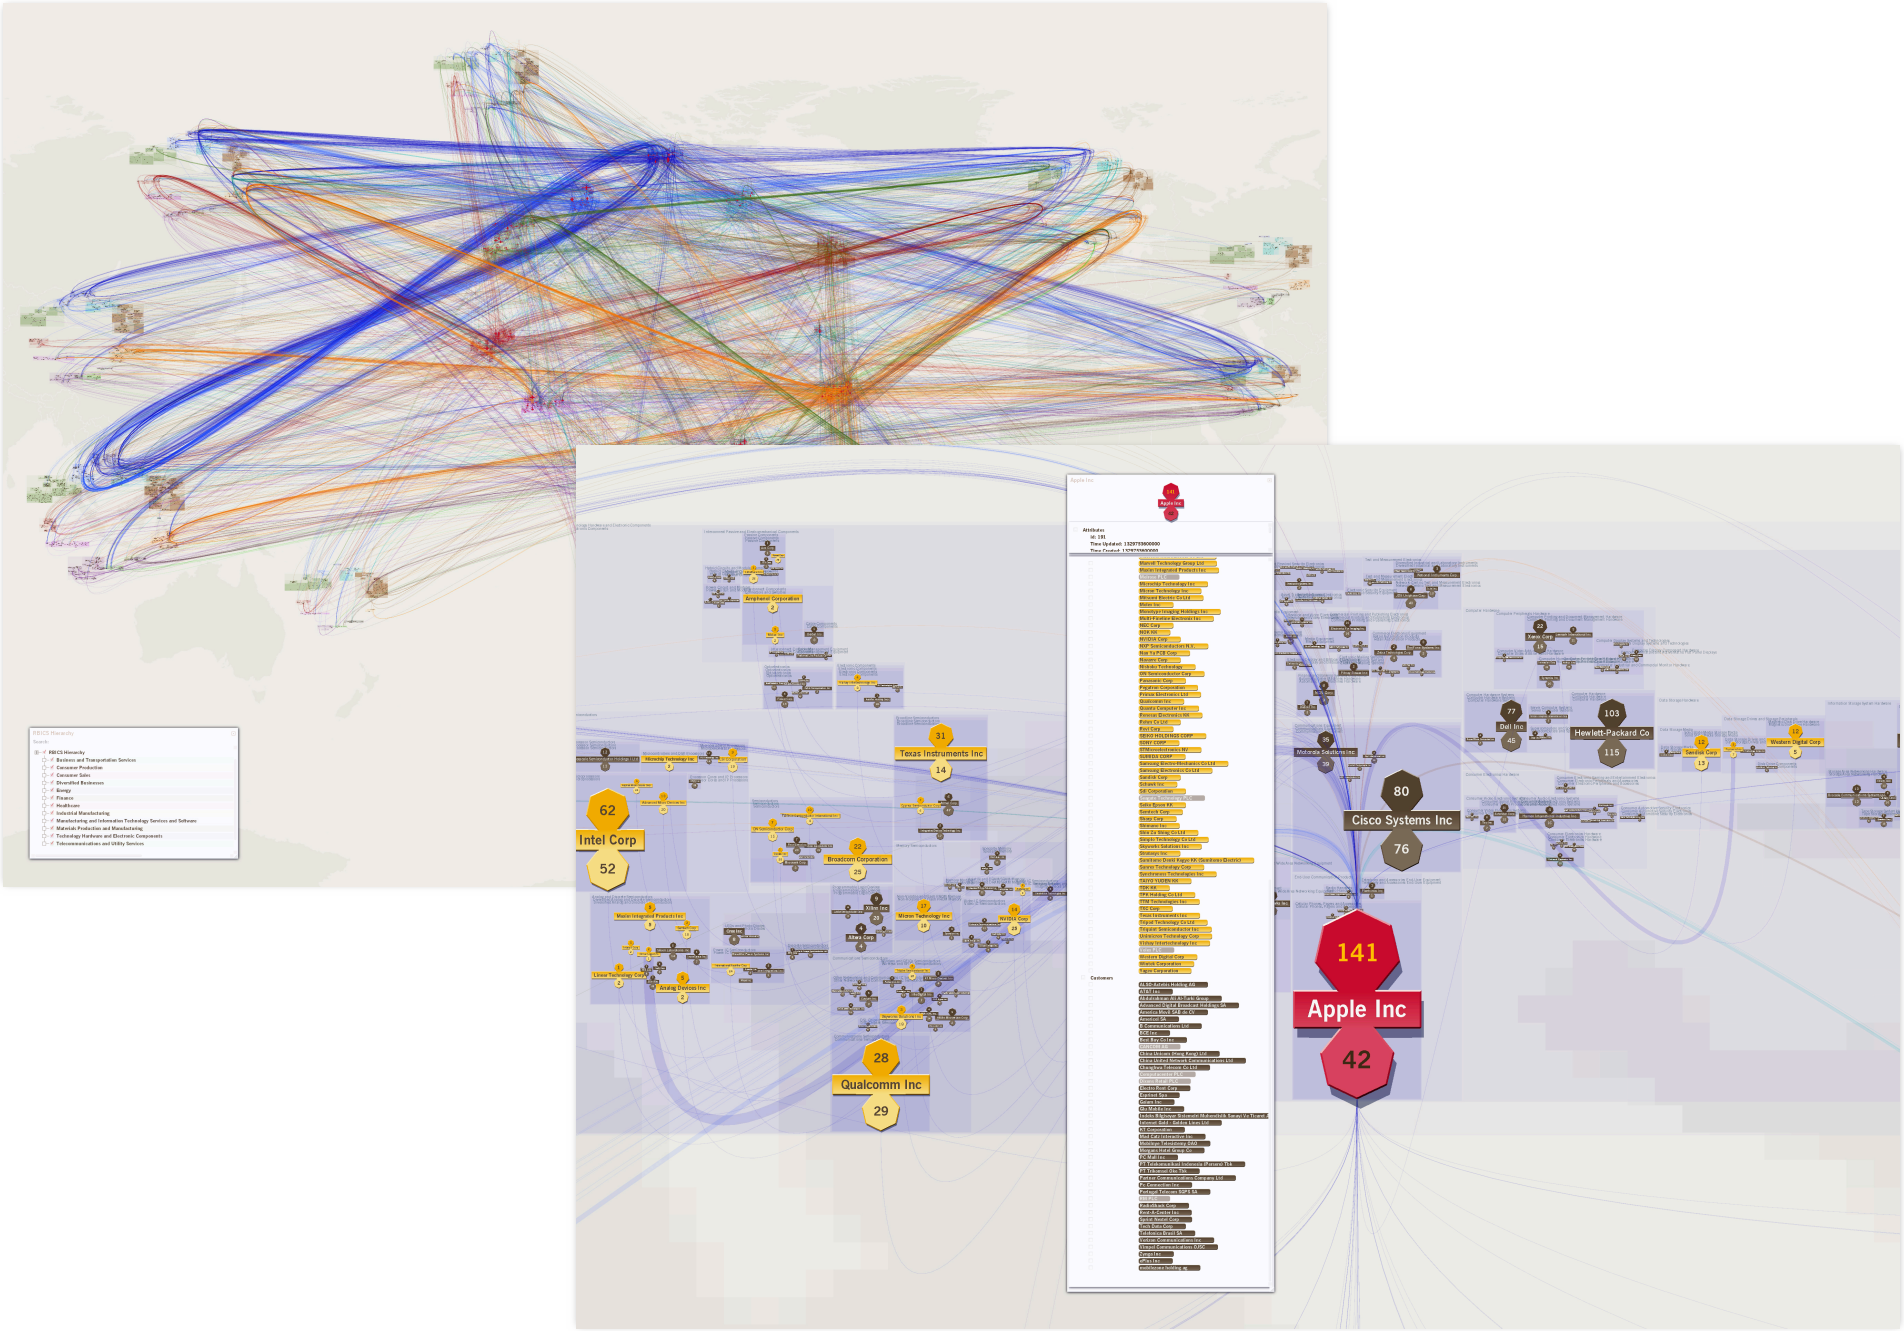
Intel Corp
62
52
Texas Instruments Inc
31
14
Broadcom Corporation
29
Qualcomm Inc
28
29
Cisco Systems Inc
80
76
Apple Inc
141
42
Hewlett-Packard Co
103
115
Dell Inc
77
45

Comparisons and What-Ifs

"Statistics is the art of with respect to what.*" A number just isn't enough on its own. Good visuals show clear, deep background, so one can judge data with respect to context.*

I worked briefly/deeply with Edward Tufte, and this simple phrase was a powerful gift. (Though we agreed to disagree on his rule of minimizing the data/ink ratio. *Ink*—his proxy for what confuses on the page—isn't the scarce resource, *attention* is. We try to minimize "data/*think*": afford more data per unit of thought effort, pardon the rhyming wordplay.)

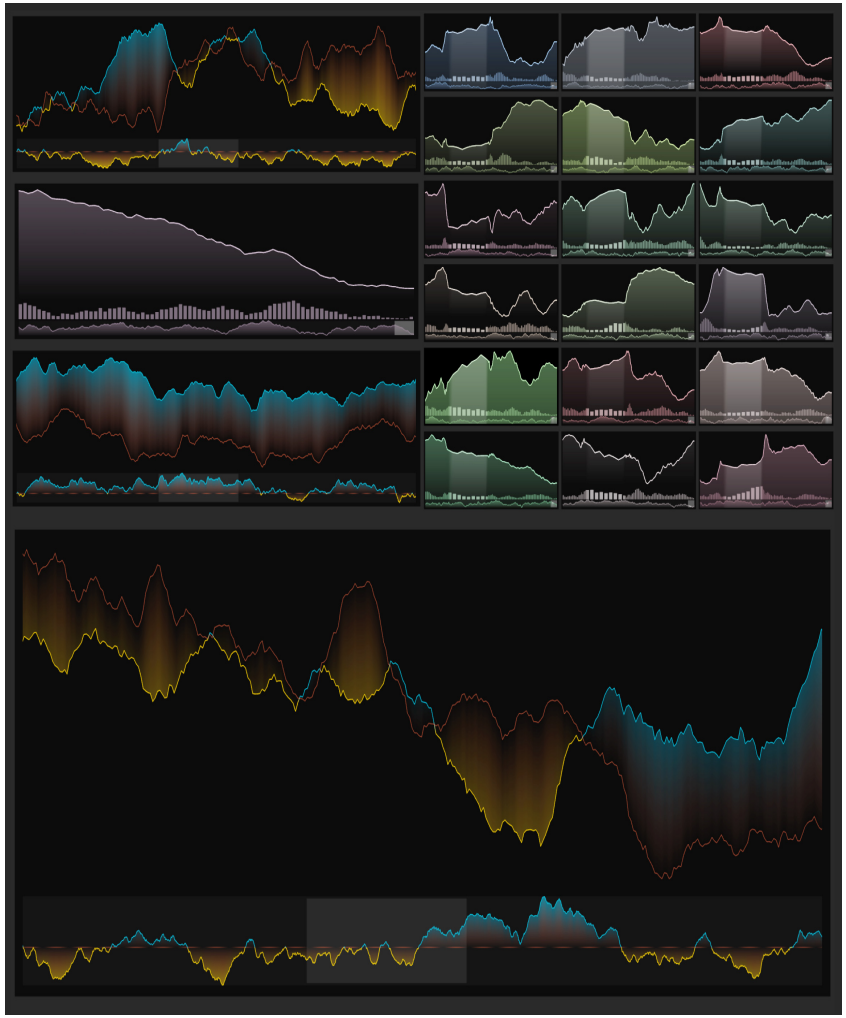

Human eyes are biased to measure gaps perpendicular to lines; it was important as we evolved to judge the strength of a branch or an enemy's arm. But it's a serious error to do that when reading charts, where only the *strictly vertical* extent of the gap encodes the difference. *Left*: We make larger gaps brighter in our charts, striations define the gap and draw the eye.

Below: Why decide based on just averages when you can see *full distributions* everywhere along the curve?

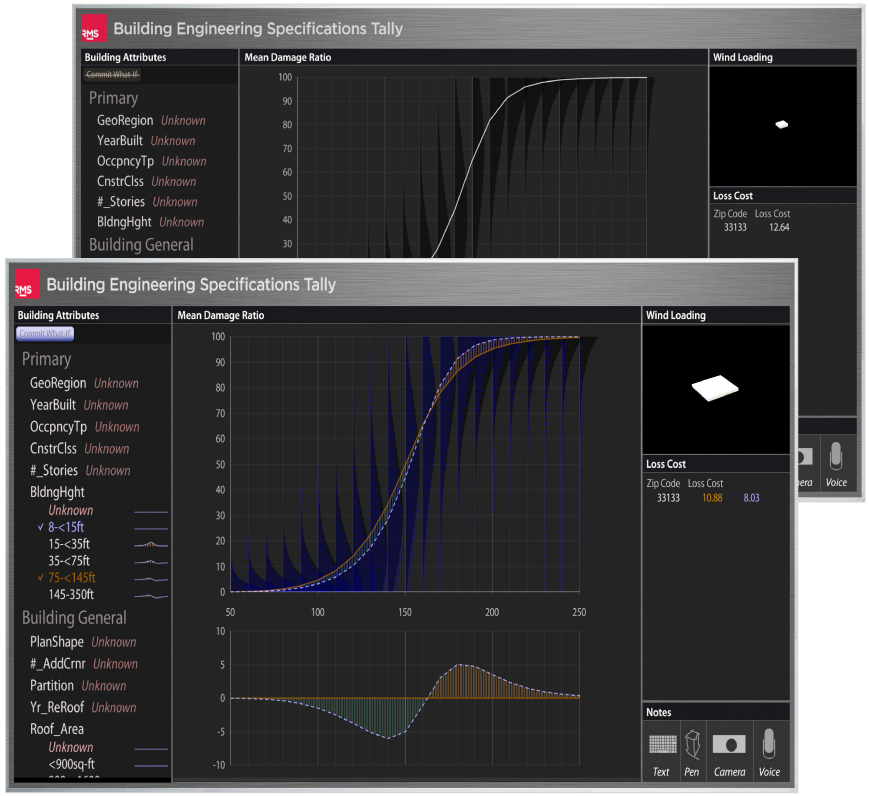

Left: It's hard to compare close lines—until we subtract them and fill the difference with color underneath. We retain the original curves above because their original shape *and* the difference both matter in this task, as in most we've seen.

Right top: contract what-ifs—in their proper intellectual context. These on-screen insurance contracts were designed so the they feel like the original real-world paper contracts. "Rapid categorization," in the perception/cognition literature, helps us find the right tool as a pop-out search, and pre-attentively activates the mental context, making us feel more comfortable when we start using the tool.

That let us develop a different visual language so people could pre-attentively distinguish the contract from the overlay, which contained the what-ifs being explored.

Right bottom: Here, the underwriter has chosen two numbers to vary, and the system shows all possible combinations in an array of Post-it-like summaries: comparing all the key metrics in the blue alternatives to the pink original. No need to go to eleven other screens, we're seeing eleven alternatives right in the expert's intellectual business context.

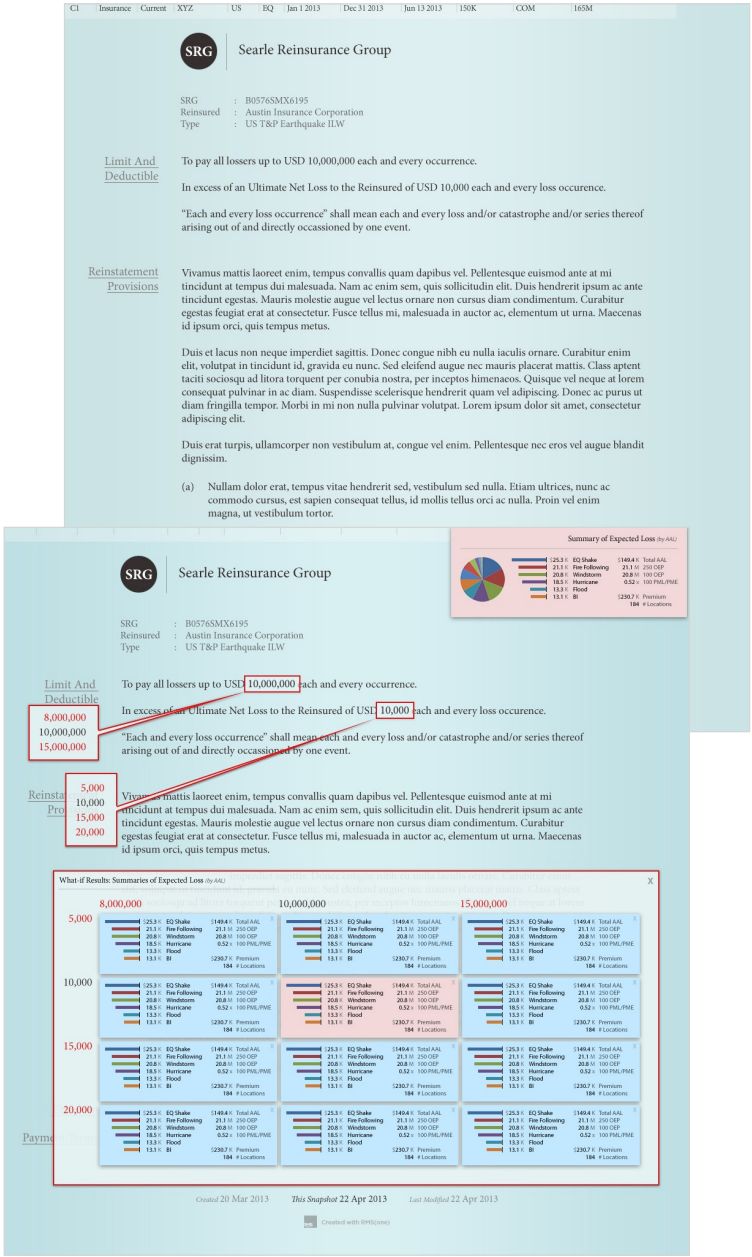

Everyday Usage

Quiet back there, please, I'm thinking…
Many visual tools use up all the available visual dynamic range, presenting themselves as if they're more important than the viewer's thoughts.

When the brief is to completely overpower a website visitor's attention and keep their eyeballs for as long as possible to convert them into a sale, it's important to use saturated, contrasty colors and tricky, thought-entraining animations. But real tools—tools glanced at hundreds of times a day—need to be humble visual backgrounds for our thought processes: clear but not demanding, structured to allow easy annotation. It's still possible to innovate with quiet, basic typography and ancient colors; just view the problem from an innovative perspective.

(Some conventions are ancient *and* exotic: "rubification" of a manuscript—adding red to foreground things—was originally done with inks made of ground rubies. We use humbler commercial process inks—but respect the visual/cognitive tradition that red means special.)

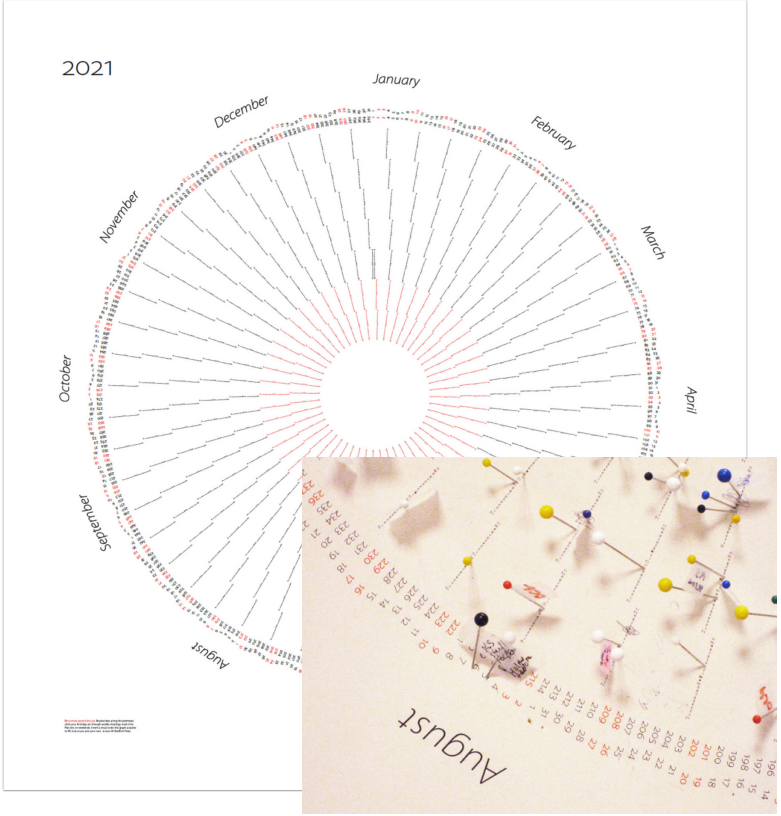

Once More Around the Sun, a year-at-a-glance calendar; every hour of the year is visible all at once, to be marked or pinned on the 2' paper version.

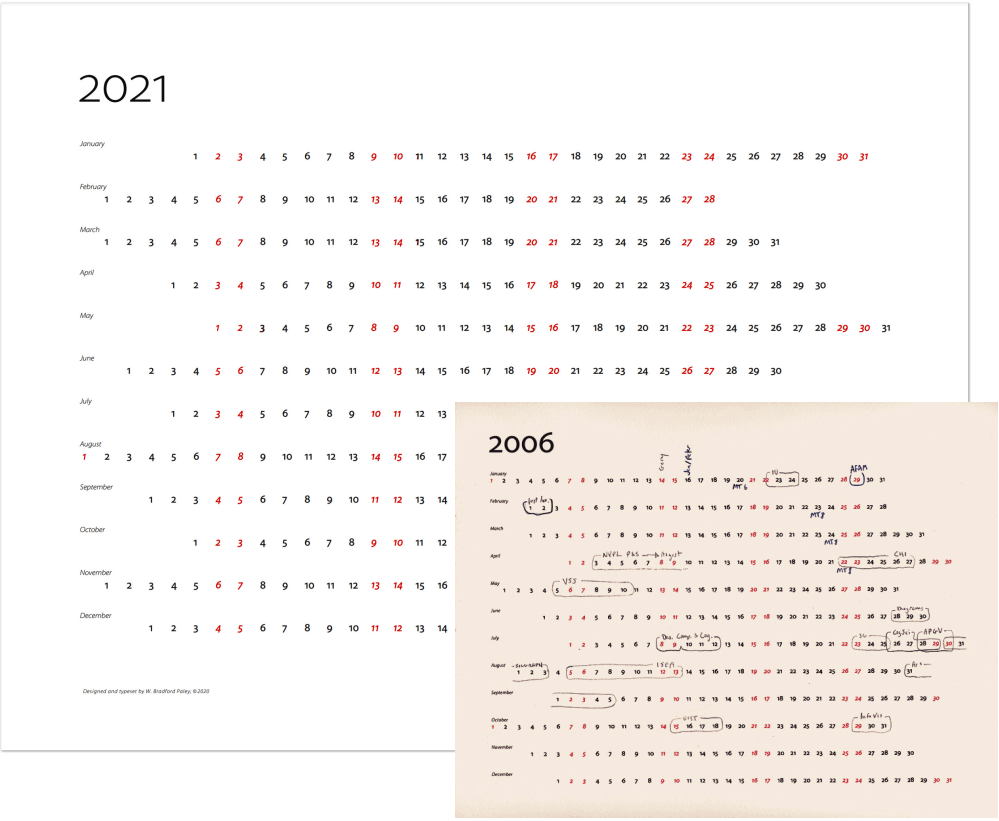

A linear calendar that breaks only at months to allow one to more easily see and arrange blocks of time throughout a year. (Here, conferences/talks in 2006; a couple had to go, and a bit of time was needed for work in between…)

Agency

Are you immersed in your task—or coping with some damned tool? A good tool acts as an extension of your body, of your mind. It can let you react and prepare before you even know an issue exists.

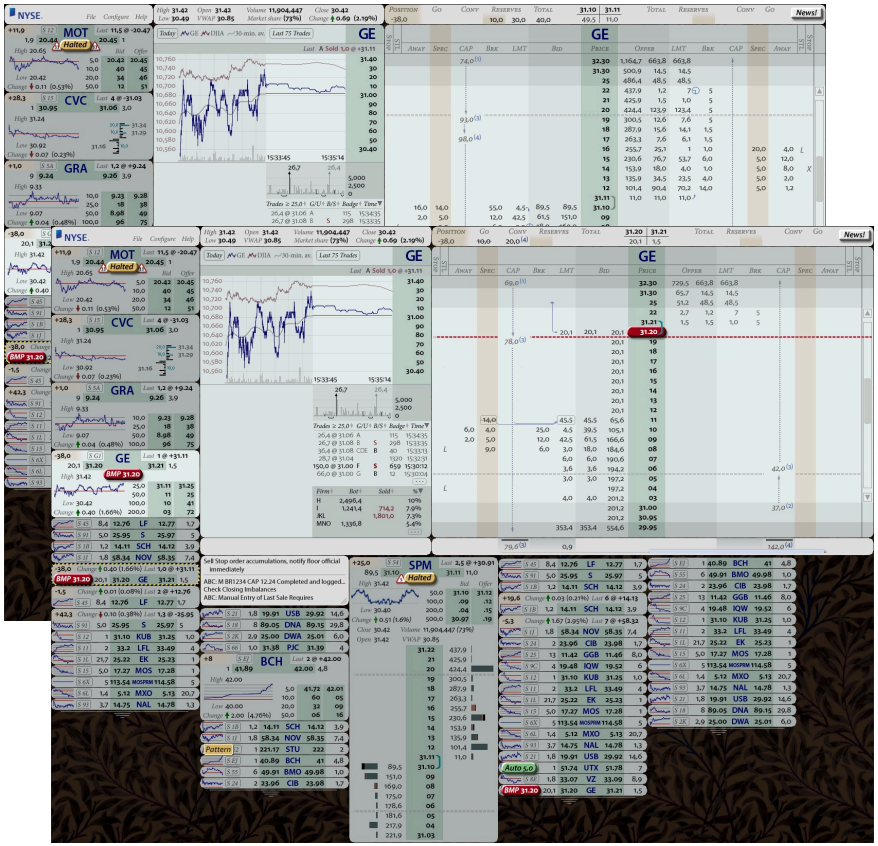

When the market's operating as usual (top) it's appropriately boring, so the specialist's eye is drawn to in-context alerts: drop-shadowed lozenges; rendered trompe l'oeil on a flat ground so they activate both our *what* and action-oriented *where* visual pathways. But if an order sweeps up through an LRP (gray dashed line, a nickel above and below the market) it's serious—things are moving too fast for market participants to post more trades in that range. That also becomes an alert, impossible to miss at its red limit line, signaling that the system has frozen automatic trading until the lozenge is clicked—signaling so clearly one can react as if it's a physical threat.

Martial artists have known for centuries that it's possible to react to something before you consciously know it's happening. It was ascribed, e.g., to one's chi being mystically in tune with the external world; like Star Wars' the Force. We now know it takes roughly 300 ms for an event to enter consciousness but less than 100 ms for deeper perceptual mechanisms to recognize certain kinds of things—so related concepts can be activated in that 200 ms—like mental scripts of how the things act and how to respond. Tools that need conscious processing, like the words in rectangles of almost all common analytic tools, don't tap into this as well as tools that present perceptual features our older visual/spatial mechanisms can process.

When tools *do* feed our sensory/perceptual apparatus properly we have a sense of agency and control. We feel like we're *working directly in the world*, just doing our task—not operating a tool. Tool *behavior*, not just visual design, is critical to functionality. People copied our GS NYSE handheld (pages 23, 28)—but just its look. Their cargo-cult knock-off got off the ground as well as stick-and-grass planes do; it lasted a single day on the trading floor.

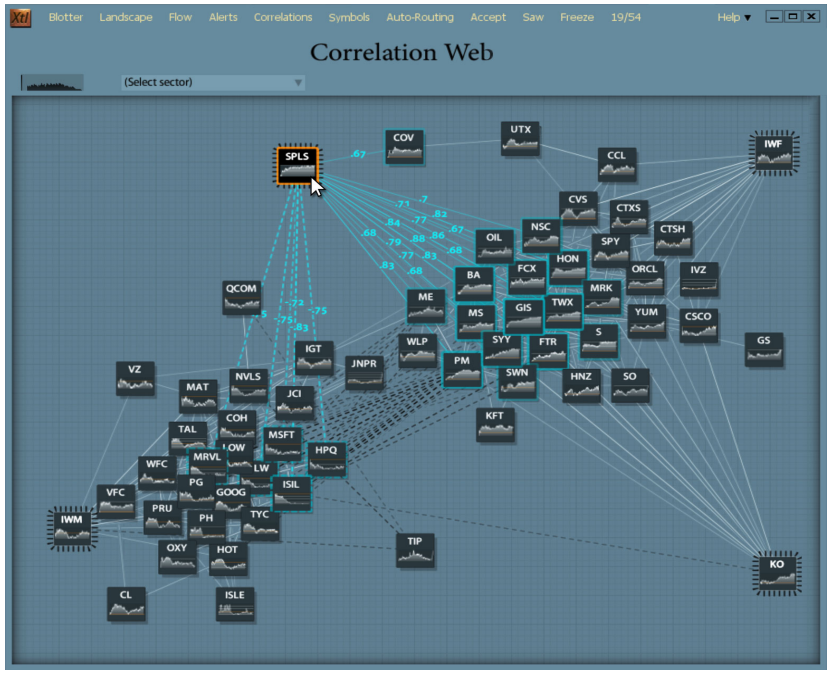

Pulling stocks connected by rubberband-like correlations creates an ad-hoc multifactor space—where the factors are defined by the nature of the stocks themselves—not arbitrary metrics. See how stocks arrange themselves, trending toward the (stitched) exemplars that were pulled apart and stuck to the background to create an ad-hoc continuum along which other related stock arrange themselves.

Flow

We feel better—and think better—when we're "in the zone." Evolution tuned us to see and react in tenths of seconds or less, so a refresh that takes even a quarter-second hobbles thought.

Mihaly Csikszentmihalyi studies the psychology of optimal experience. But when you're "in the zone" your whole-mind behavior goes deeper than the veneer of experience. In fact, flow's pleasant experience may be a by-product of working effectively in the world: an evolutionary reward making it more likely to be selected into the next generation.

Csikszentmihalyi lists conditions that induce flow, and we've developed many techniques that create those conditions. Notably: a tool that looks like a specific task reinforces a clear goal, and its affordances' direct sense of agency create the clear, immediate feedback a good tool provides: Ben Shneiderman's "direct manipulation" with "rapid, incremental, reversible" operations.

When a Web tool takes time to respond it's far worse than a lost ½ second—it interrupts flow, which has been tested to sometimes take more than 23 minutes to regain. This ruins productivity and alienates the viewer. The Web might someday grow out of its "terrible twos," interrupting us because it feeds on our attention, or because pages are poorly implemented. But a mature set of tools avoids this already, independent of the platform.

Flow is hard to illustrate on a static page, but here we see *direct manipulation*, a flow-engaging interaction paradigm where one directly manipulates the representation of the data.

We see it in three forms: *first*, in the Attributes view we're sweeping out the range of data we want directly on the graph, directly on the data.

This creates three data "outports" which, *second*, let us directly/tangibly pipe the "within"-range partition of the data into a *SectorCircle's* "mark" inport.

By removing the need for button pressing or writing code we essentially let the person's *body* do the data handling, offloading that effort to older mental mechanisms where task-appropriate and freeing the mind for analysis, insights, and collaborative work.

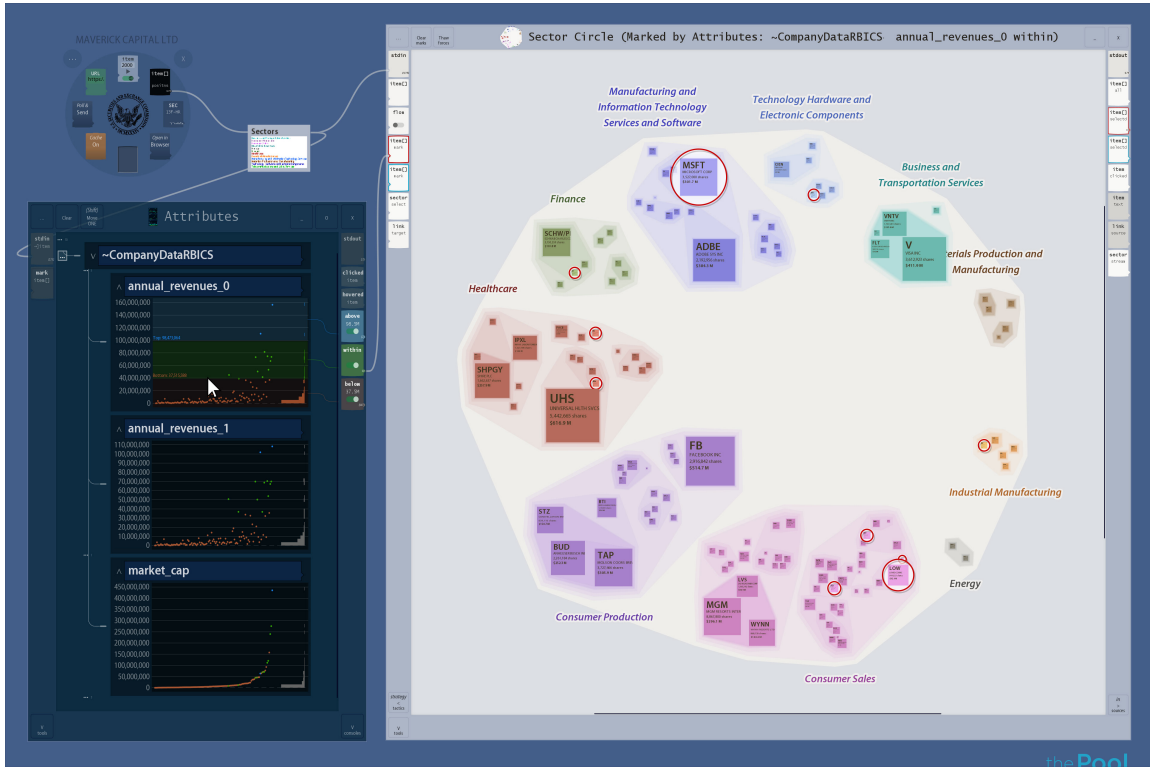

Third, to mark more positions we can pull the range bottom up an down and instantly see the marks disappear and proliferate until we've marked (circled) just the number we're looking for.

No iterative query necessary.

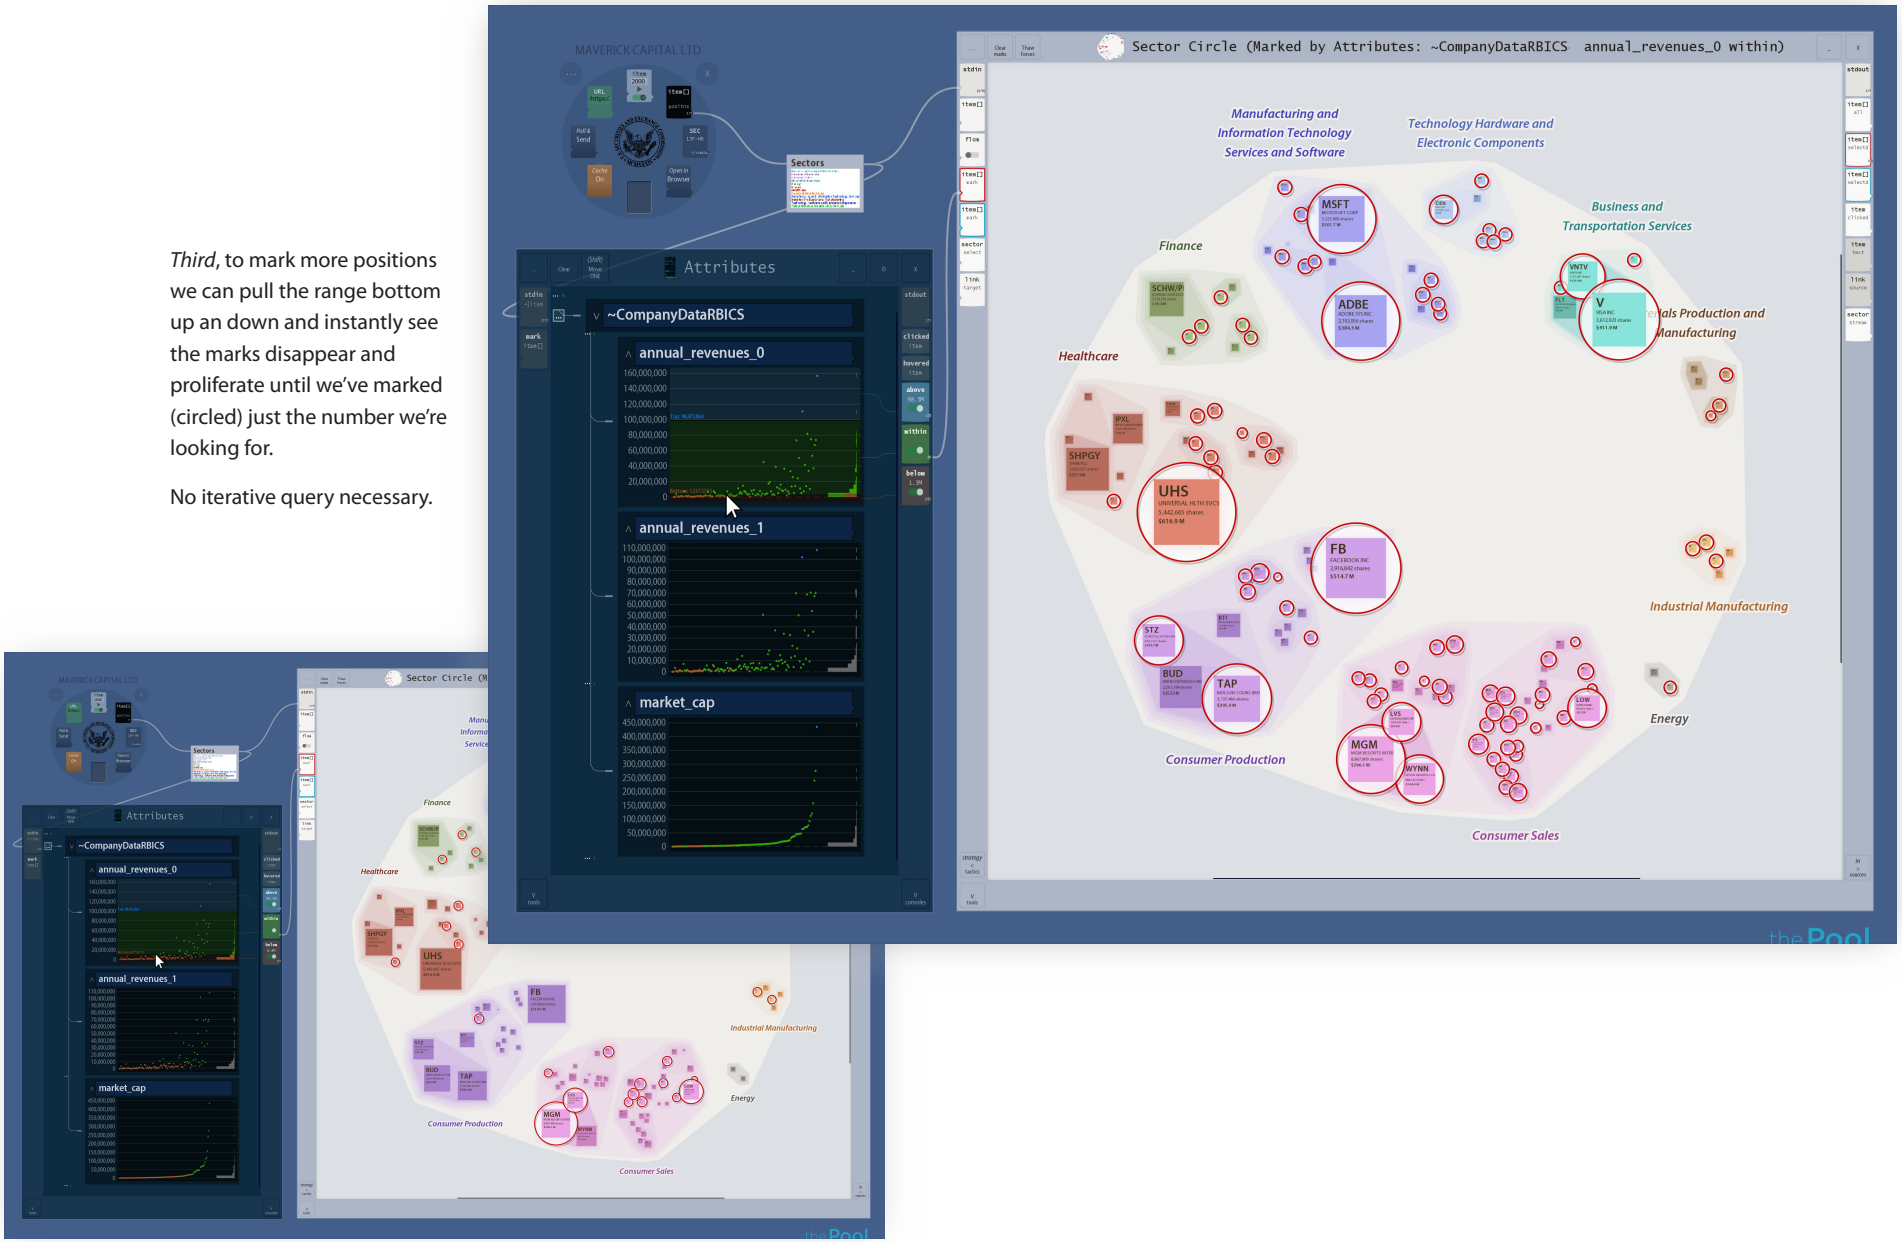

Chunking & Functional Encoding

Manage Visual Complexity

We "rescued" the traders at a major international energy-trading bank from a design created by a salient global markets UX-design firm. Their design followed all the proper design rules; our followed the way traders think.

Making a separate field for each input certainly created a tidy-looking strip that needed to constantly be scrolled to complete the task; one of the primary complaints.

And worse: It. Broke. The. Thoughts. Of. The. Traders. Into. Many. Separate. Pieces. Worse. Than. This. Typographic. Trick. Does.

Doing our discovery right on the trading desk we just listened. Those pieces were used in *sentences*—each one a meaningful and known subtask: a cognitive chunk in their workflow. So we created these Composite Fields, which cemented the separate fields back together into those real-world chunks.

An unsolicited email from the project director reported: "feedback has been positive, people booking complex products without having to understand the application or been given training."

No scrolling at all. And in startup-speak: 100% customer satisfaction, zero training costs, zero support costs. All done within a flat design vocabulary.

Facilitating mental "chunks" lets people juggle more.

Former UX Design

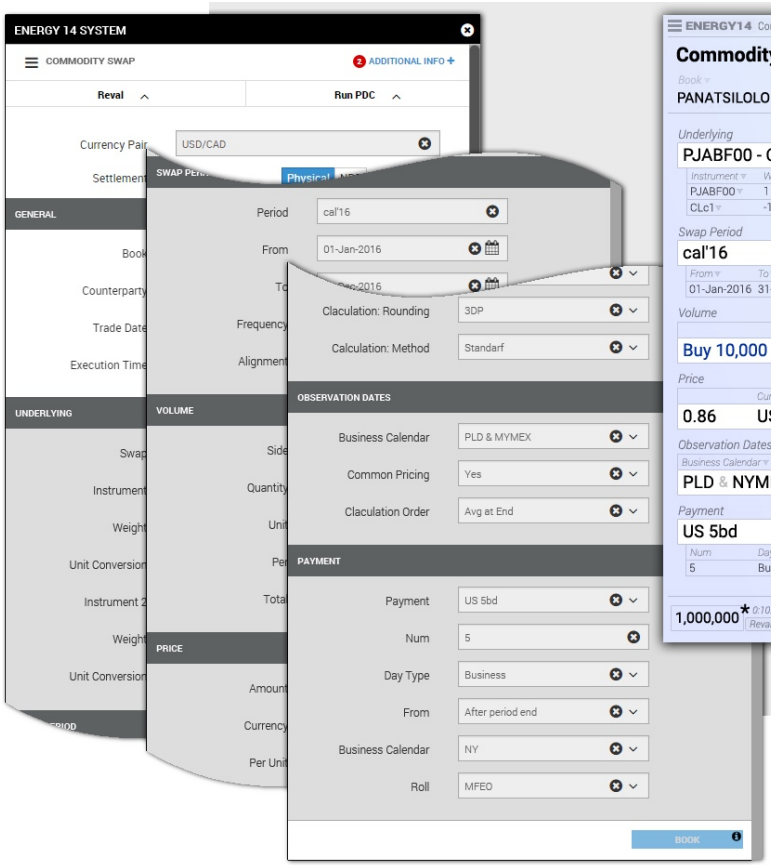

Our Redesign

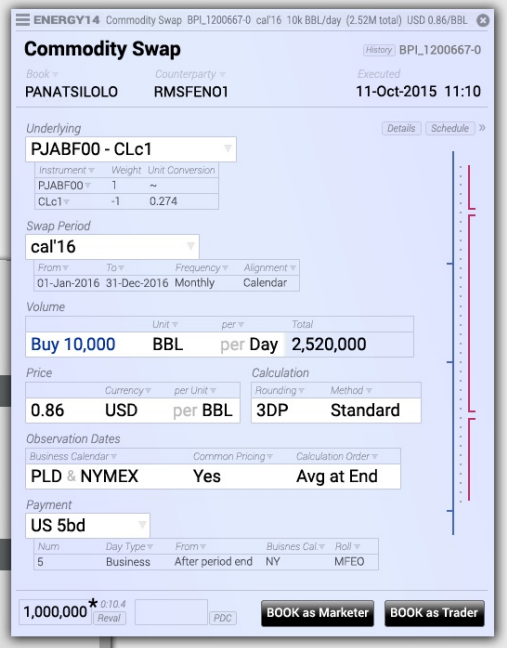

This gray background is the target screen size; the previous UX team's design exceeds it by 3x.

Our redesign fits the target screen easily and allows more recognition information, e.g., a blue/red linear payment schedule graph.

Ironically, we often work to shed *attention rather than attract it; visual differentiation helps us ignore what we're not looking for.*

Another bit of history, the Goldman Sachs broker handheld from 1999, has strongly-differentiated visuals for different functions so your eye snaps to what you need—by letting you ignore the things you don't, just like in the real world.

Flat-UXer fashionista denigration of skeuomorphism misses an essential point: we animals *need* those rich visual attributes to encode real-world objects by their functions. Imagine one of them redesigns your home with even their best UX widgets for every function—and there's a fire. Do you survive?

Left: The ink-entry screen copies that era's floor-broker physical paper execution slip (*inset*) because it worked beautifully—it evolved over decades to an optimal size & configuration. *Center:* The reality-mimicking keypad can't be mistaken for anything else. *Right:* When brokers needed to compare several orders to choose among, since copies of the paper ones didn't fit on the screen we transcribed their *meaning*, not the ticket.

Here's a cognitive-engineering almost-koan: *how do you illustrate something someone's never seen, yet let them recognize it instantly?* This one has an answer: entities are just bundles of features, whether they're in the mind or drawn. So if you draw the idea's characteristic features, encoding importance and relationships, they're recognized and mapped back to the mental entity—voilà: people *see* their thoughts.

Visual features in this design instantly show all the actionable parts of an order: what's buy or sell, what's amended/working/done. And critically: organizing by price let brokers ignore all "out of market" orders, focusing where deals happen: sells touch buys "at the market." The previous design hid that, ordering by time so they could find new ones at the top. Here new ones are in-situ—unmissable black until they're touched/accepted.

This design kept the Goldman brokers the fastest on the floor, competitive for years after NYSE electronic trading started.

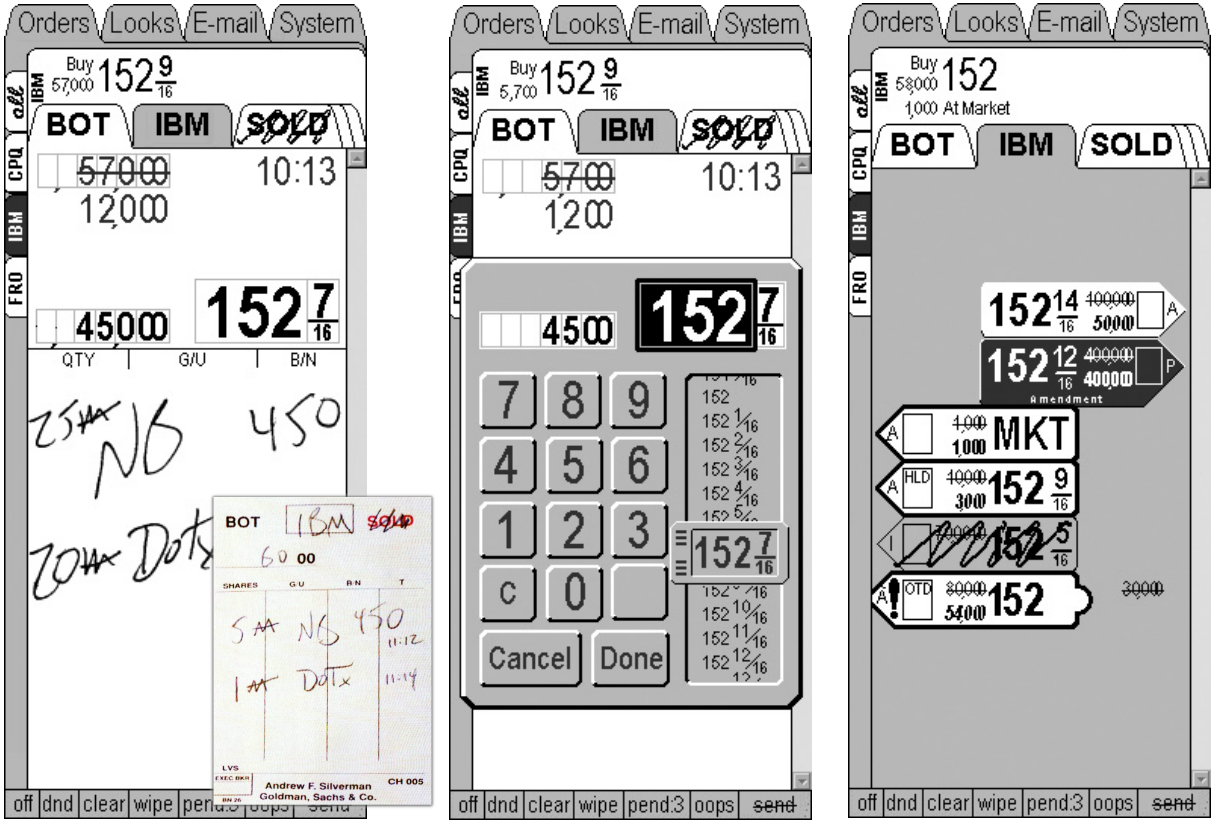

Co-workspaces & Presentation

Hosting distributed cognition. Peers think with *us when visuals act as shared memory; the ideas right there in front of people in other roles.*

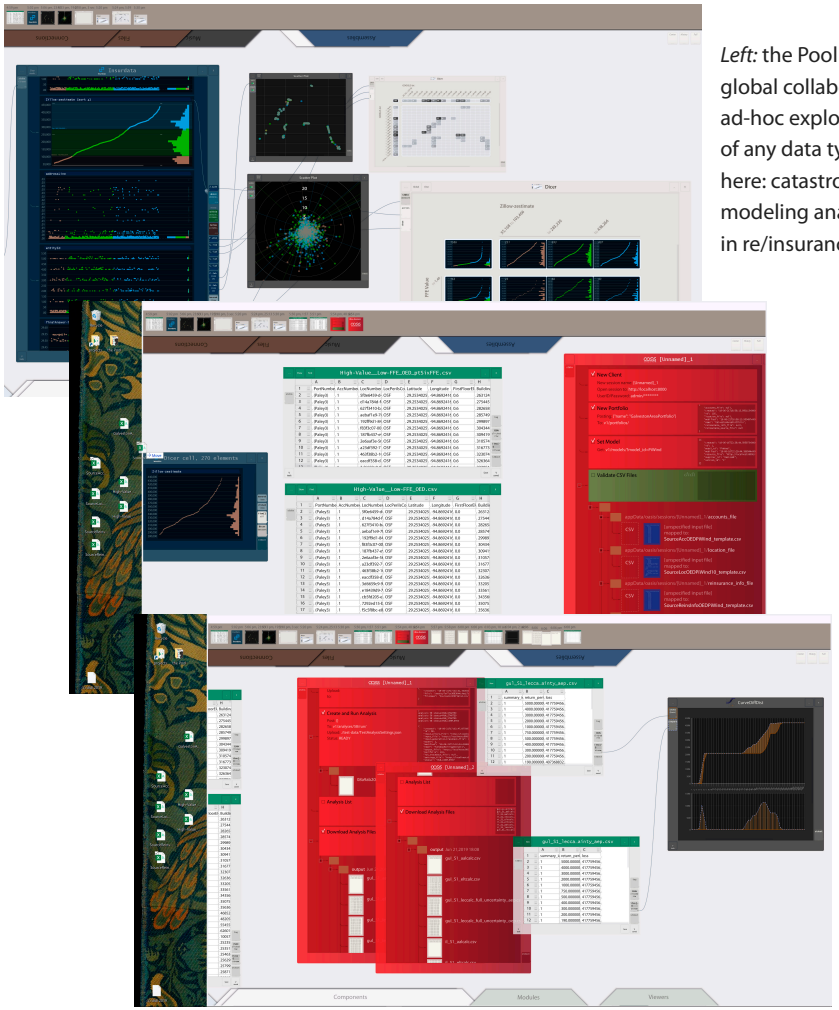

Left: the Pool allows global collaborative ad-hoc exploration of any data types; here: catastrophe modeling analysis in re/insurance.

Edwin Hutchins studied how the five Navy roles on a ship's bridge become a single cognitive entity. Their tools (map, compass, log book, even the physical throttle position) bind their thoughts into a single emergent intelligence, able to do cognitive operations no individual could alone. Entities that emerge like this need not be self-conscious like science fiction's hive minds to work, and have existed as long as humans—in countless forms from tribes to quartets to partnerships. People in closely-interacting roles, like analyst/portfolio-manager, or PM/investor *could* be such entities but they're divided, collaborative thought all but prevented by today's isolated tools.

Our tools enable these meta-minds to form by enabling instant flow of information between the people involved. They act as shared memory, facilitate non-verbal communication—and optimize group-think.

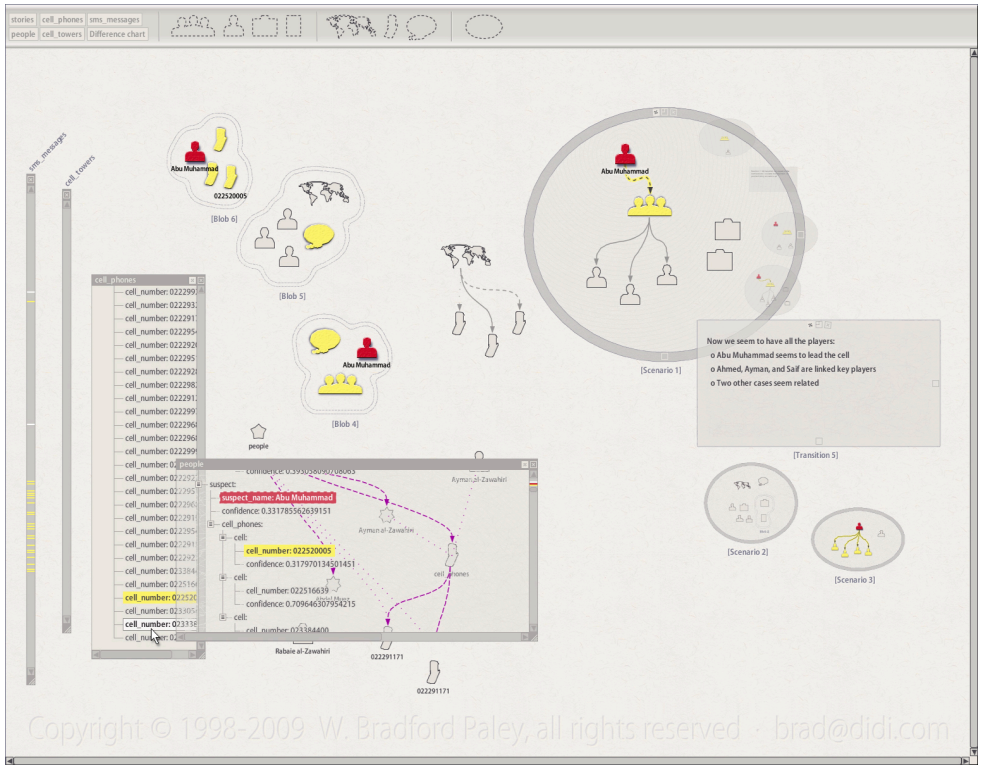

Above: WorldMerge let intelligence analysts with different backgrounds, skill sets, even working in different agencies, work together to synthesize quantitative data with qualitative data—whether they were in the same room using the same display or across the world.

Right: Brad Paley is interviewed at the Whitney Museum in front of his work *CodeProfiles*, which reveals coding and optimization processes to a lay audience.

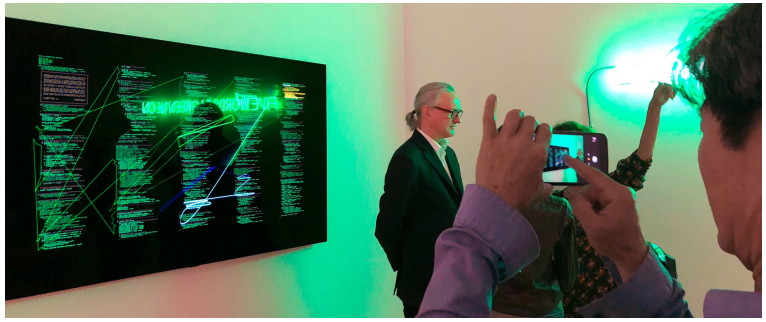

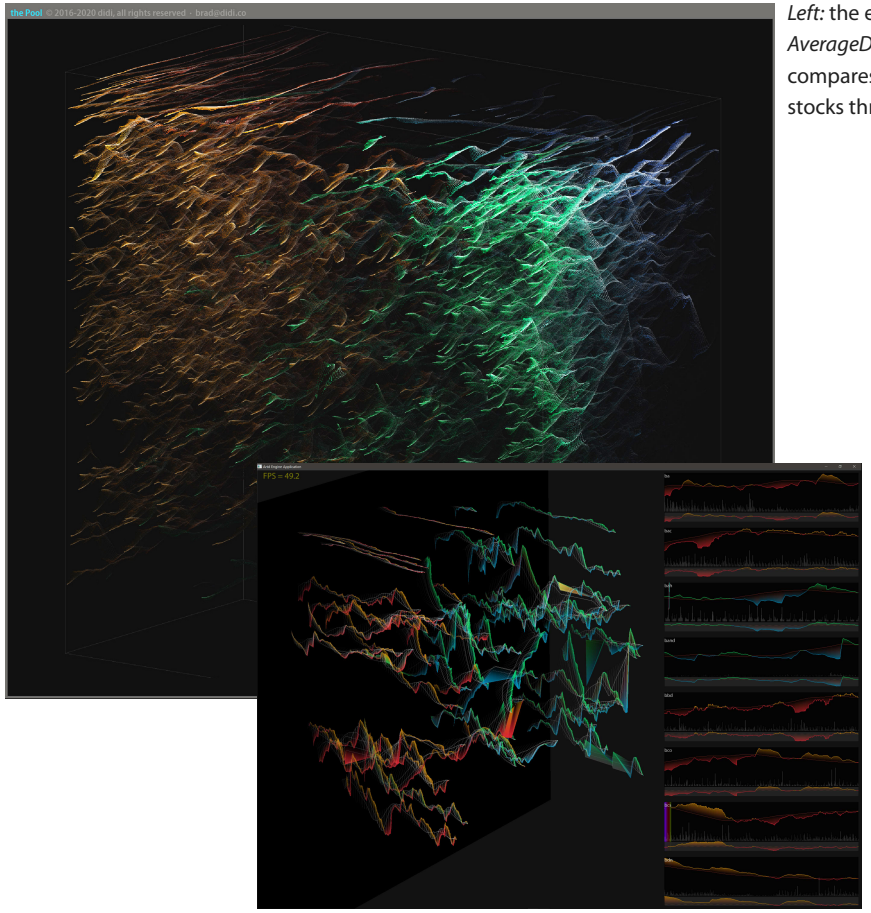

Left: the experimental *AverageDragons* view compares hundreds of stocks through time.

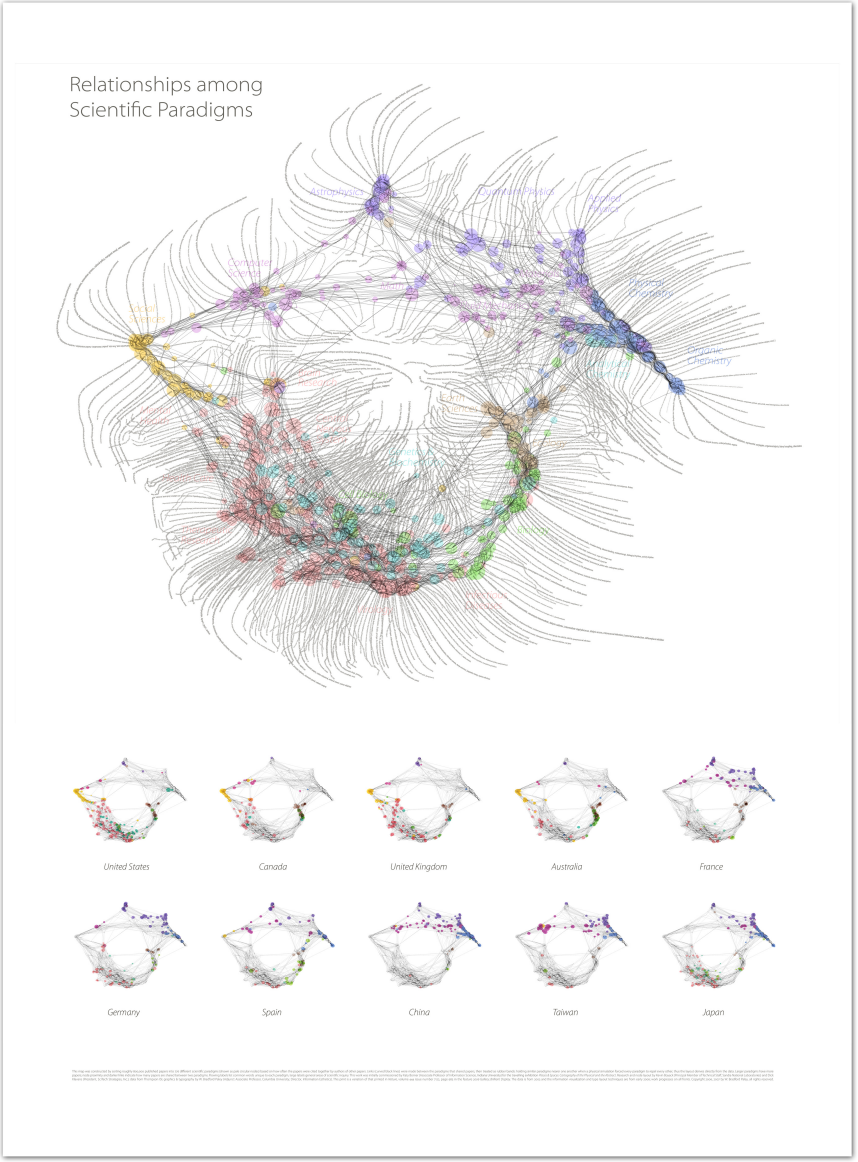

Above: a map of science (featured in *Nature* and a dozen other publications) organizes the whole world's scientific publications. The ghost of that large map is then repeated below for several countries, coloring just that country's contributions. The original was three feet wide, but look closely and you may see the West's strength in medicine (*the map's lower left*), the East's strength in math and physics (*upper right*), and Europe's relative balance.

2 Applied to Real-World Tasks

A sampling of products/tools designed and/or implemented by Digital Image Design Incorporated

The Process

The process we employed in creating these tools delivers results that are more effective than those coming from common product development processes

If these images seem more meaningful and engaging than most analytics-support systems, if they look richer, more clearly grouped and organized, or if they more immediately telegraph domain-specific concepts and actions, it can be attributed to Cognitive Engineering. Much of this work is as close as we can get to directly transcribing domain-specific concepts and actions. This is what we think of as true, unconstrained Cognitive Engineering Design.

But even the more prosaic-looking work embodies a subset of Cognitive Engineering we call UX++. UX because it operates in a browser, engaging the mental apparatus people develop using common widgets; a designer's usual excuse for not inventing new engagement models.

Created back in the 60s and 70s, these widgets *do* fit form-based tasks—which still exist. But the intervening half-century has applied little more than cosmetic make-overs, leaving an aching need for at least two significant increments. One + results from each task action being analyzed rigorously by experts *in the task's domain*, as well as experts in tech and design—generally at the granularity of sub-second cognitive operations. The other + foregrounds a willingness to include subtle visual elements that directly and uniquely depict specific business entities: pre-attentively recognized; zero cognitive load to find and use. *In*consistency serves recognition.

Every color, margin, line, font, tint, transparency, shape, and drop shadow in this work was tuned for visual grace on the page—*after* it was shaped to guide the eye more quickly and surely to what's important for the task, to evoke insights more easily, to inform decisions to be more rigorously balanced, to afford more fluidly-executed actions and more easily-confirmed results.

You are invited to fan quickly through this whole book to absorb this work as a gestalt, noticing how different we made the tools look in order to correctly reflect the variety of tasks supported. But you're also invited to scrutinize tiny details, to try to unpack the meaning those details depict, even if you're not an expert in that domain.

Better: absorb these techniques & visuals as an expanded palette we can together use to illustrate your *own* domain. The process of making a painting is regularized, but the results vary widely, revealing each subject's uniqueness. This is how to create interfaces that don't *interfere* with the task, but disappear as if they're *windows into* it.

Deep visual semantics of task details differentiate this work; they stabilize fleeting concepts, from the mind to the screen (or paper!) so people can build second-order ideas on top, so they can enter flow state more effortlessly. These "pre-understood" visual representations of domain elements *get the system out of the way*. People look not at the surface of the screen but *through* it—directly recognizing and applying their ideas within the thought space of the domain to which they've devoted their lives.

Cognitive science now recognizes that even abstract thoughts largely have a foundation in physical metaphor, grounded in space and shapes. Articulating the interface to leverage the spatial mechanisms of thought increases productivity, reduces misinterpretation risk—and most important: facilitates insight, creativity, and the personal satisfaction experts chose their domain for in the first place.

NYSE Floor Broker Handheld

Goldman Sachs, 1998

Key feature: Invention of the "information object" representation to reduce cognitive load reading orders (*near right*). It reassembles order attributes into an illustration of the order thought-entity actually living and working in brokers' minds. (Attributes are usually torn apart and spread across a grid as its columns, they still—suboptimally—are in almost all of today' common transaction/analytic tools.)

"It's hard to let go of something you've been doing for a long time and do something new. Brad's given us an ergonomic shorthand. You don't need to learn it, you just grab it with your eyes. The bottom line is that it makes the broker more productive. And that makes the firm more productive."

—Andy Silverman, Goldman Sachs broker, in *ID Magazine's* 4-page feature

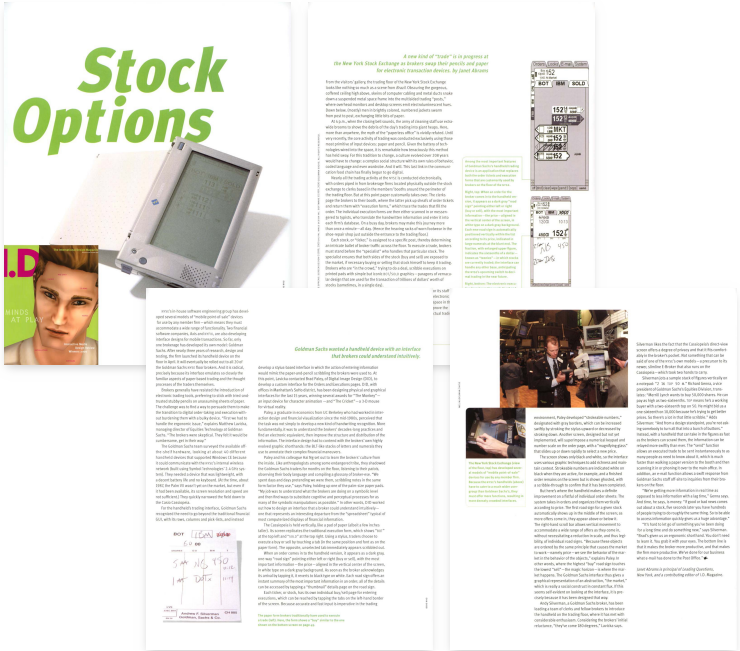

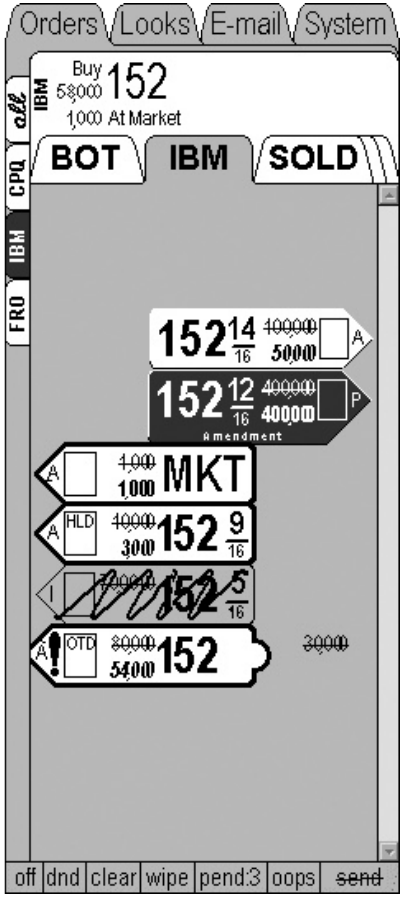

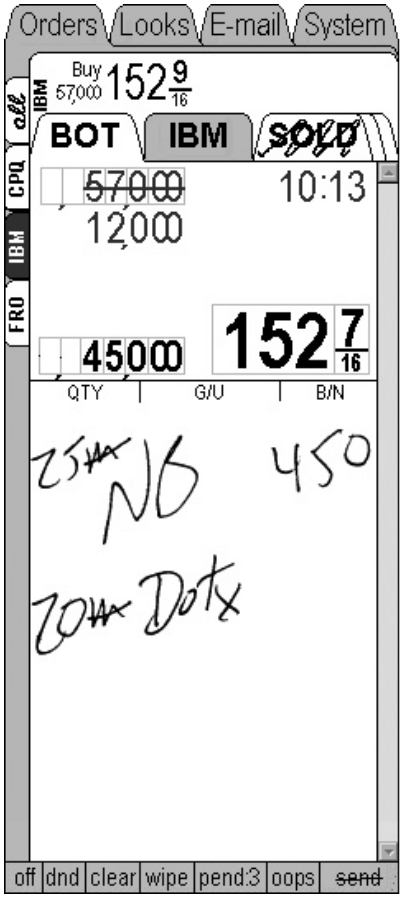

Equities Trading System

Merrill Lynch, 2000

Key behavioral advantage: Created clear visual trade tracking within the sales trading blotter: every trade's "fills" progress and state was immediately visible, as were "unmissable" alerting visuals for critical operations.

Accessibility advantage: Conventional (red/amber/green) colors were used for alerts, but they were tuned to have the proper differentiation and visual saliency ordering for color blind people, 7-8% of all men.

Development cost savings: This is an early application of the Paper Prototyping stage in our design methodology. By illustrating and printing out every significant system state, then hand-simulating the system in front of a trader asked to pretend they were using it, we caught bugs immediately and added enhancements before both coding *and* even spec-writing—this saved huge amounts of effort, money, business-side frustration, and calendar time.

"I worked with Brad at Goldman Sachs and thought his jet-cockpit-like ergonomic approach was normal—until I tried to find a team in London to do the same. We looked at everyone we could and didn't find a fraction of the background, insight, or capabilities. I was happy to be able to use didi once I got back to the States. The resulting interface clearly made our Equities Sales Trading more productive as well as happier using the new system."

—Perry Metviner, Managing Director, Equities Trading Development

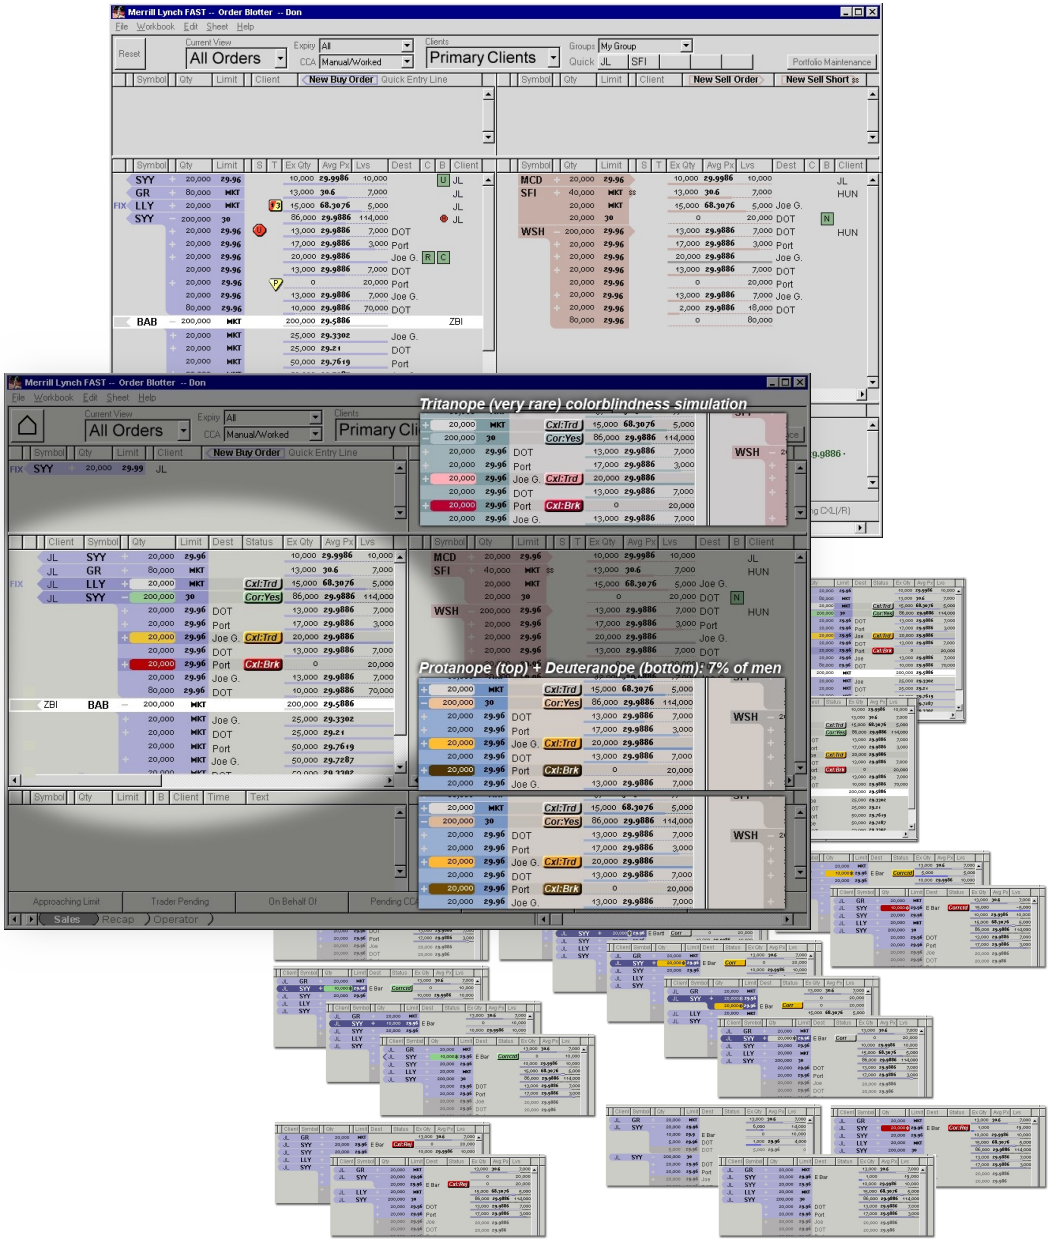

Tool for Structuralist Text Analysis

TextArc, 2000

TextArc is a graphical index of a text that can be used with any text. It generalizes to simultaneously show frequency of use and distribution of any set of entities with respect to a linear ordering of another class of entities that relate to them.

TextArc draws every line of an entire text in a tiny font, wrapped into a clock-like ellipse inscribed within the display bounds. This alone results in an unreadable mess. To make it readable, it also:

- Draws each significant word at the centroid of its usages in the text; pulling it close to where it's used most
- Draws those words larger and brighter if they're used more
- Draws star-like "distribution glyphs" next to each word to show where it's used in the print
- Draws rays from words to where they're used when hovered in the interactive program
- Draws each word in a color related to its position: high saturation means a word is used mainly in one place, hue is determined by its sequential/angular position in the text; this allows words common to one area to be more easily seen because the color contrasts with the context when used elsewhere

TextArc has been exhibited in The Museum of Modern Art, the Chelsea Art Museum, Google's New York headquarters gallery; it won Grand Prize [non-interactive] at the 2002 [6th] Japan Media Arts Festival; and had many other appearances "in person," on the Web, and in books.

Patented in 2003: US 20030235807

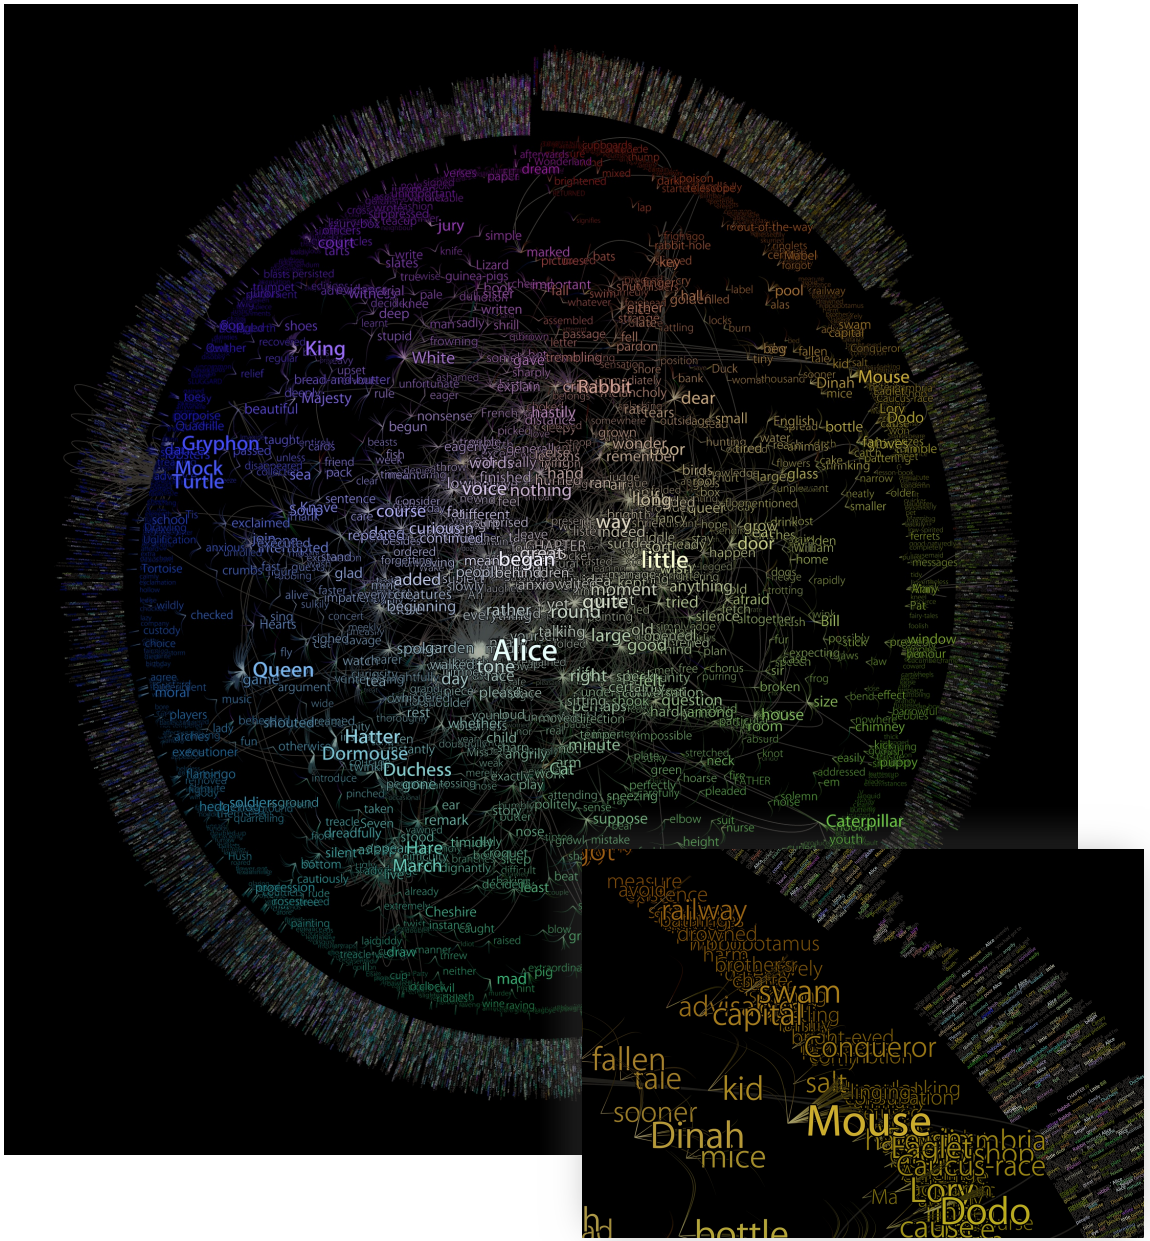

Foreign Exchange (both sales & trading)

Morgan Stanley 2002

Key advantage: Created "unmissable" dimensional lozenges as alerting visuals (*right*) within the trading blotters—dangerous situations were made impossible to ignore by visual layering.

Separately, on the quick input screen (*below*), the risk of trading with a client at a deficit was reduced by introducing a real-time market feed that put real prices in the execution buttons and made prices that lost money look dangerously red. Note that the trader could still do the deal; important if they'd verbally committed. Often overly-zealous business analysts design limits into the system that get in the way of unusual but sometimes-necessary business practices. We *alert*, not prevent. We trust the expert to know their job.

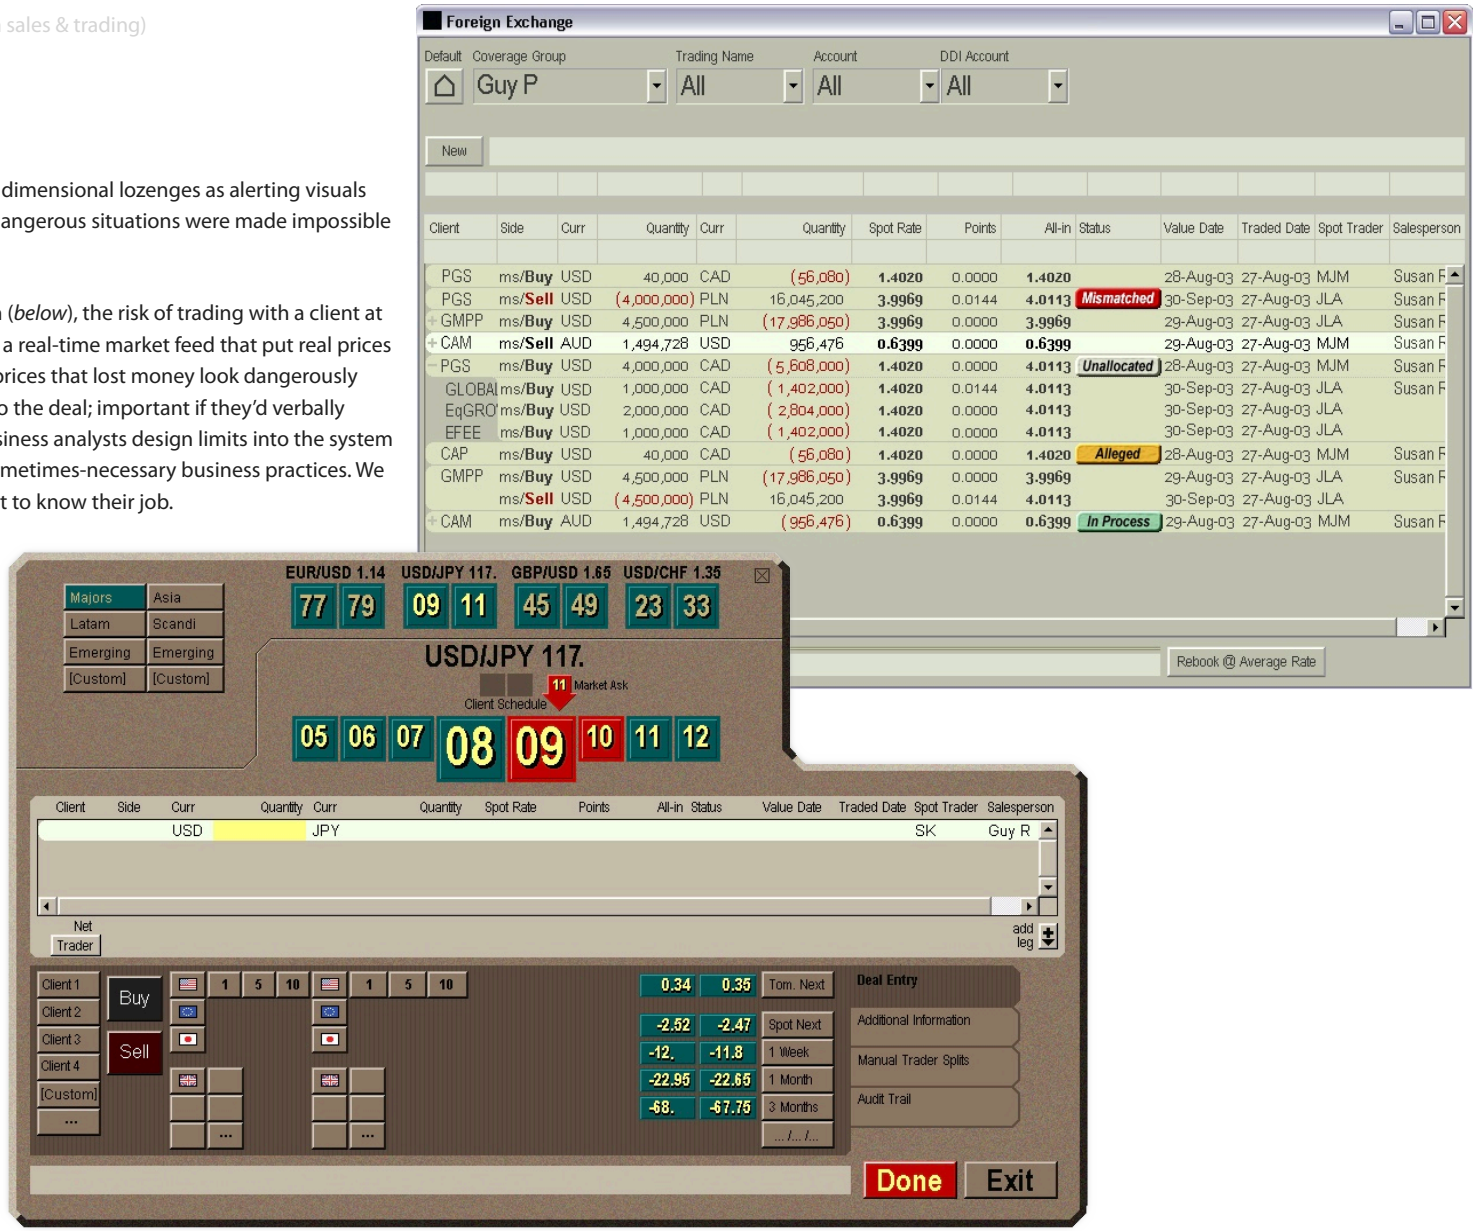

Content-Guided Search & Scan

Kellogs, 2003

CoGSS was an early search tool designed for medium-sized document libraries. It was tuned for library exploration and senior management use. Long before Large Language Model AI, it tried to assist a search effort by looking through partial results and gently suggest them to the researcher. A conceptual design sketch is to the right.

The layout is an easy-to-use left-to-right refinement search, with four panes:

- Common Terms: the search tool suggests what to look for
- Search Criteria: multiple ways to look for a document
- Document List: the list of documents that meet the criteria
- Document Preview: a curated excerpt of a document

Each pane has special features:

The Common Terms pane may have three lists in it, each one showing many possible search terms (e.g. Health, Atkins Diet, Organic, February 2003).

The three lists are currently "Favorites" (a list that people can populate themselves), "Business-Wide Terms" (terms that commonly occur within the business), and "Terms in Selected Documents."

This last list is rather special: it looks through all of the documents that are in the current "winnowed down" list of documents and shows the terms most commonly used in those documents. It can be a help for people to see what sorts of topics remain, and provide "proven" (known existing) terms to refine the search. It changes every time the Document List changes.

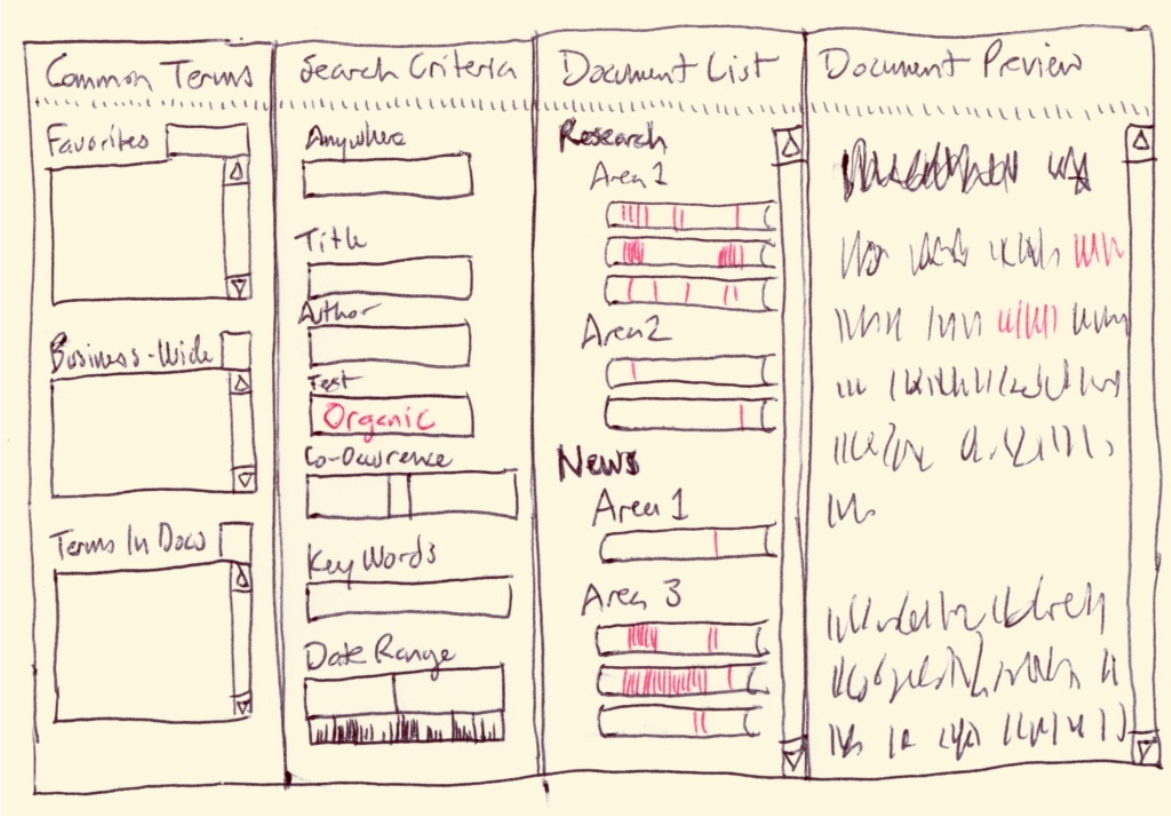

Specialist Workstation

New York Stock Exchange, 2004

Key advantage: Put up to four times as much information in the same screen space, if information is measured by characters on the screen—yet left it more readable than before; put perhaps *fifteen* times as much information on the screen if measured as inferences that the specialist can make.

"Brad Paley looks at the world differently, and provides a perspective in design that is innovative and intuitive for users. He is able to combine his unusual perspective with years of proven successful design and implementation to deliver transformational change with the mastery of a minimalist."

—Lou Pastina, Executive Vice President

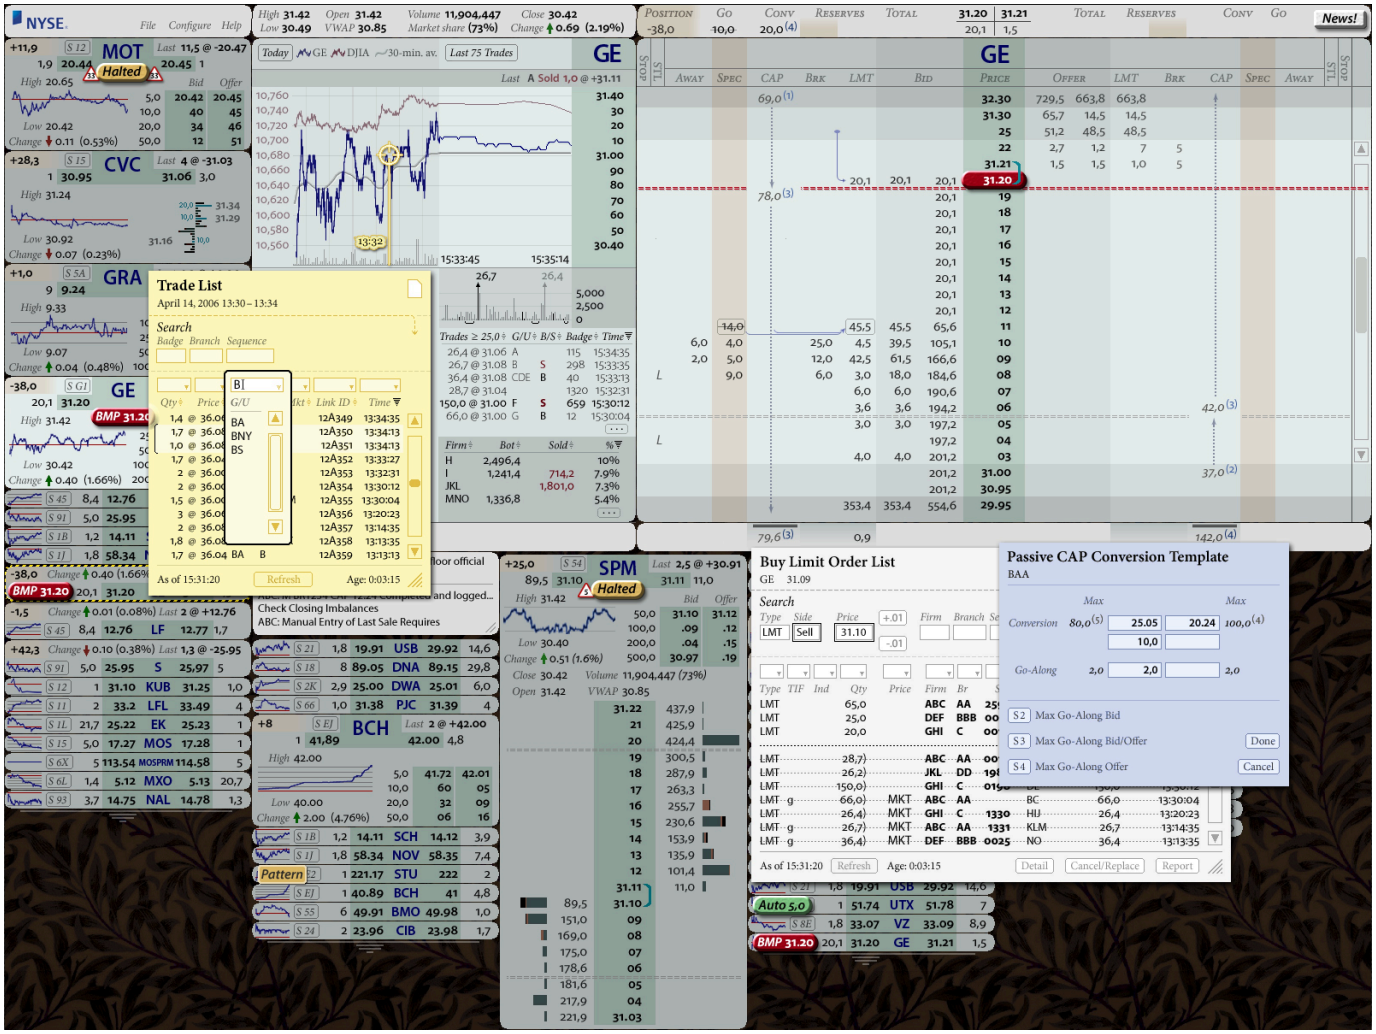

Orion

Morgan Stanley, 2005

Orion is a sophisticated internal organization chart, HR database, and access permissioning system.

Key transformation: Extremely abstract representations of complex relationships among teams, people, roles, permissions, securities, systems, clients and other entities were made concrete; easy to drag-and-drop, easy to modify.

Cognitive Engineering Principle: map data back to the real-world elements and behavior they model. We used common, recognizable screen objects to evoke real-world entities (e.g., a person's outline for a person, a briefcase to hold the data needed for a specific job). Then we carefully matched the expected real-world constraints implied by those objects to the system's data constraints.

For instance, it's common sense to drag a person into a group, or put 3-holed paper (signifying data type) into a binder, then to put either a binder or piece of paper into a briefcase—but not to put a person in a briefcase. We took system operations out of the realm of abstract UX forms and invisible, hard-to-remember associative rules and recast them as expected, obvious relationships among things; this reduced errors and simplified learning.

Operational Impact: 50+ Web-based screens which could previously be operated only by the program developers were reduced to six objects & three views. The resulting system was easily used by the entire intended audience throughout the firm.

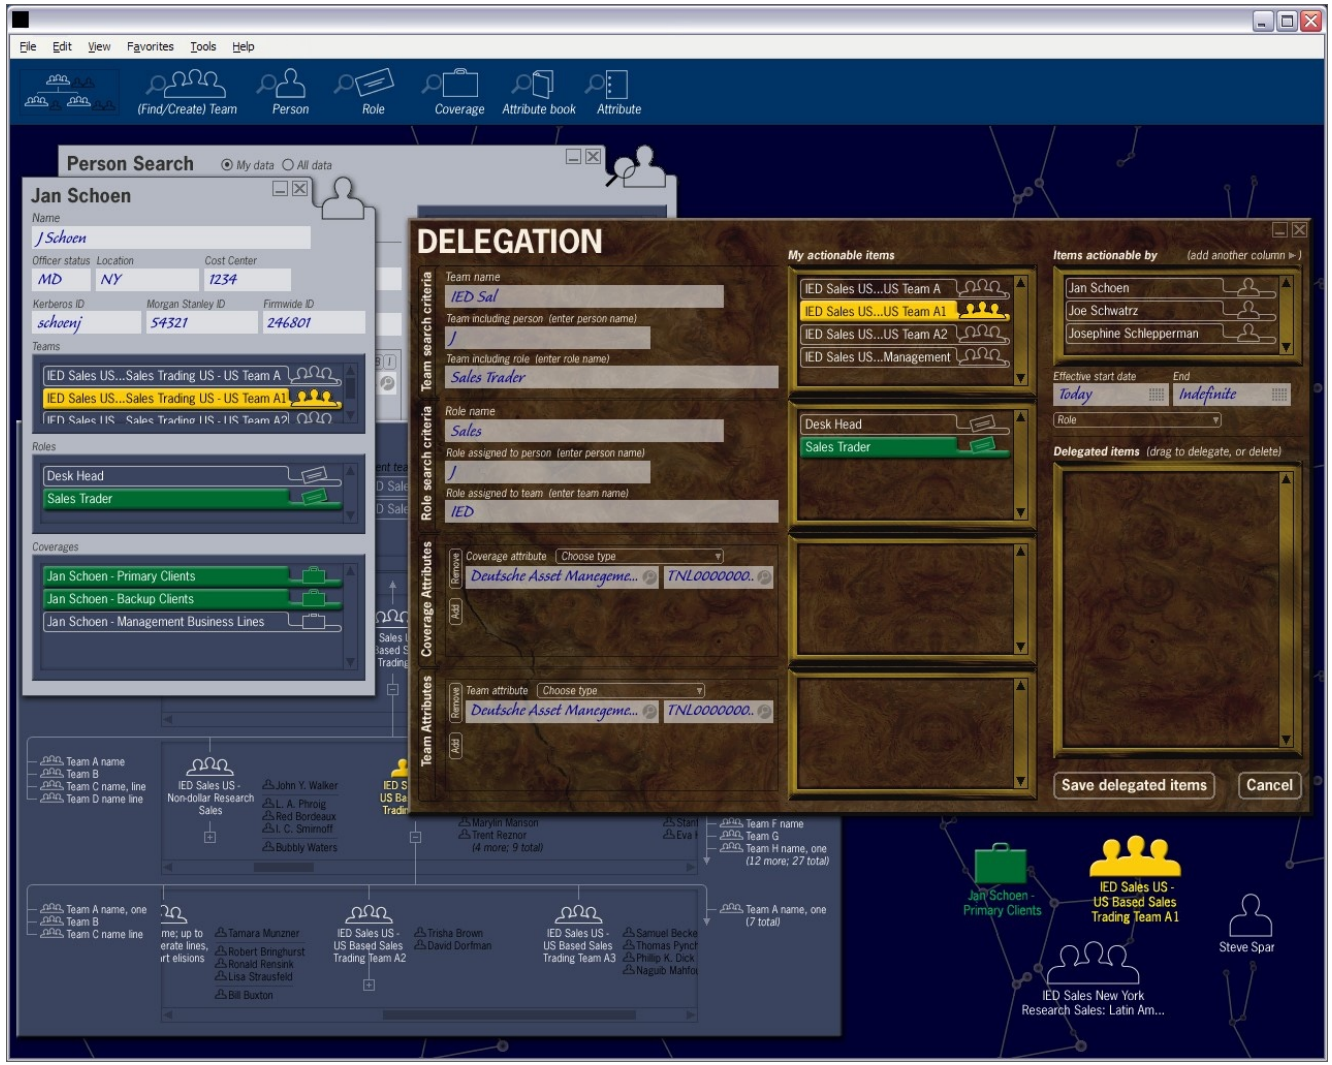

Broker Handheld

New York Stock Exchange, 2006

Key advantage: Sped up brokers a factor of fifteen to twenty (7-10 seconds was independently tested to reduce to 0.5 seconds) in the fundamental process of doing a "refill" order. We also greatly sped up cancel-and-replace operations for moving one or more orders to a different price point while enabling improved strategic thinking and reducing risk: just drag the child-order tile (e.g. *5,0*[q], clicked, so outlined in blue lower left) to another price point.

Fundamental business contribution: some of these orders had non-trivial behavior. The clicked *5,0*[q], e.g., had a discretionary range of +4¢, and followed 2¢ below the MKT (where buys meet sells), hoping for a big sell order to sweep down and execute it—but had a ceiling of 26¢. Quite abstract, and previously hidden (sometimes forgotten) in a separate details screen.

We illustrated these behavioral rules (light blue overlay lines) directly in the context of the market—so they could be understood and even modified while thinking about how they related to all currently-active orders. (The up/down arrows could be dragged to change the penny boundaries of the parameters).

In fact, these were the earliest algorithmically-driven electronic trades. Though contemporary algorithmic "strat" trades are even more complex, and some play out in statistical spaces rather than market spaces, this same approach could be used to bind algorithmic bots and human flexibility into a real-time market-compensating hybrid. Chainsaws are powerful, but better when a human directs them.

"Apart from his great skill as a designer, Brad has a unique ability to develop a design by listening and understanding what clients want and need. At the NYSE, Brad was asked to design a state-of-the-art portable order management system. His constituency was a broad group of high level brokers with disparate needs. By combining his communication skills with his design knowledge, Brad developed a remarkable product."

—Gordon Charlop, President and CEO of broker WJ Dowd

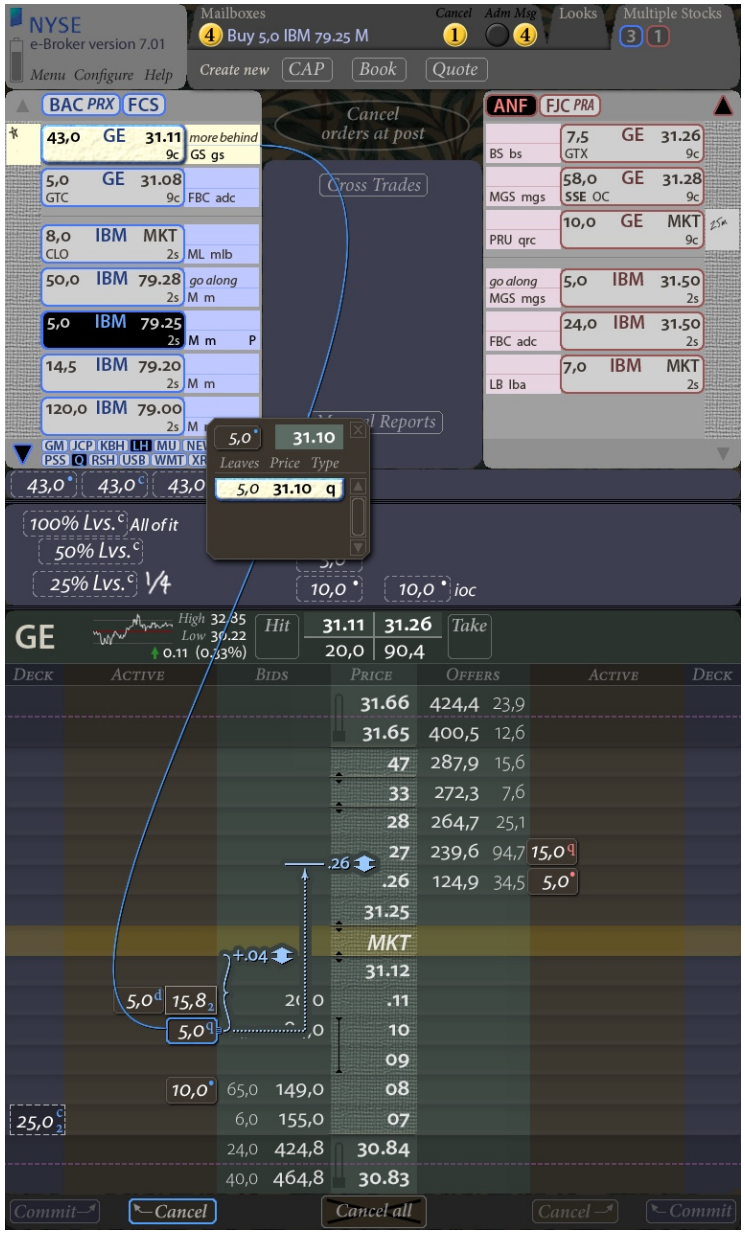

Prime Brokerage Client Report Access

Lehman Brothers, 2007

Key advantages: Vastly simplified the definition of customizable reports and specification of automated report delivery by both clients and internal staff, organized immediate access to those reports. We're told the project manager later duplicated the innovative ideas at Citibank as her own.

Illustrated here: One of the working pages, originally printed on real paper approximately three feet tall to allow easy Paper Prototyping iteration with the core team; we printed scores of these. All process discovery and sub-second cognitive task decomposition was done with a varied team of experts including coders, client support people, and system experts.

Cost & Image-Risk Benefits: Each Web page took weeks to design and hone, but once it was implemented it was released immediately to the entire audience of thousands of clients—with no need for re-design or interface bug fixes. This is a primary advantage of Paper Prototyping: virtually no post-deployment "agile" fix-up releases are ever needed—they can go toward adding new functionality.

"Our Prime Brokerage Web site was complicated, but the didi redesign made it both simpler and clearer. The didi approach made our highly-functional (but complex) client-access Reporting system much easier by applying unique didi information-object technology, and the paper prototype-based design process we learned from didi saved us uncountable development hours and business-side frustration."

—Perry Metviner, Managing Director, Equities Trading Development

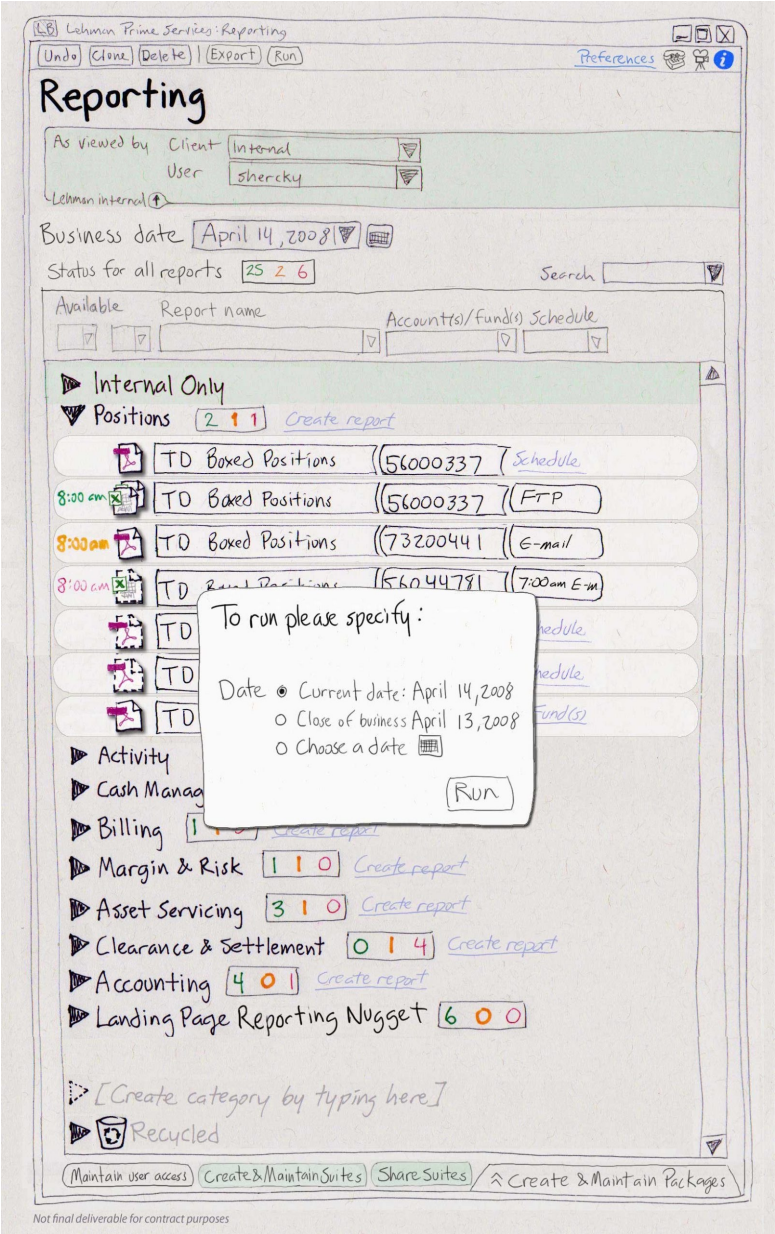

WorldMerge

Earle Industries, 2008

Intelligence Community analyst's workstation prototype.

Key innovations: Newly innovated interaction techniques allow instant "data fusion" association of disparate classes of data objects in freeform "blobs," and relationships between objects with categorized links; the design allowed several analysts to work together, in the same room or globally, in real-time or asynchronously; this is our earliest work explicitly innovating toward Distributed Cognition. The concept of two-sorted multiple logical analysis was directly visually supported; "scenarios" (smaller ellipses) organized information related to a developing "story"; and different stages of the formation of a story are represented in an active visual audit trail.

"We came to Brad because common sense in cognitive flow is apparently rare. We needed someone to guide some conceptual thinking very early on in the project and have found his approach invaluable."

—Ted Goranson, Project Director

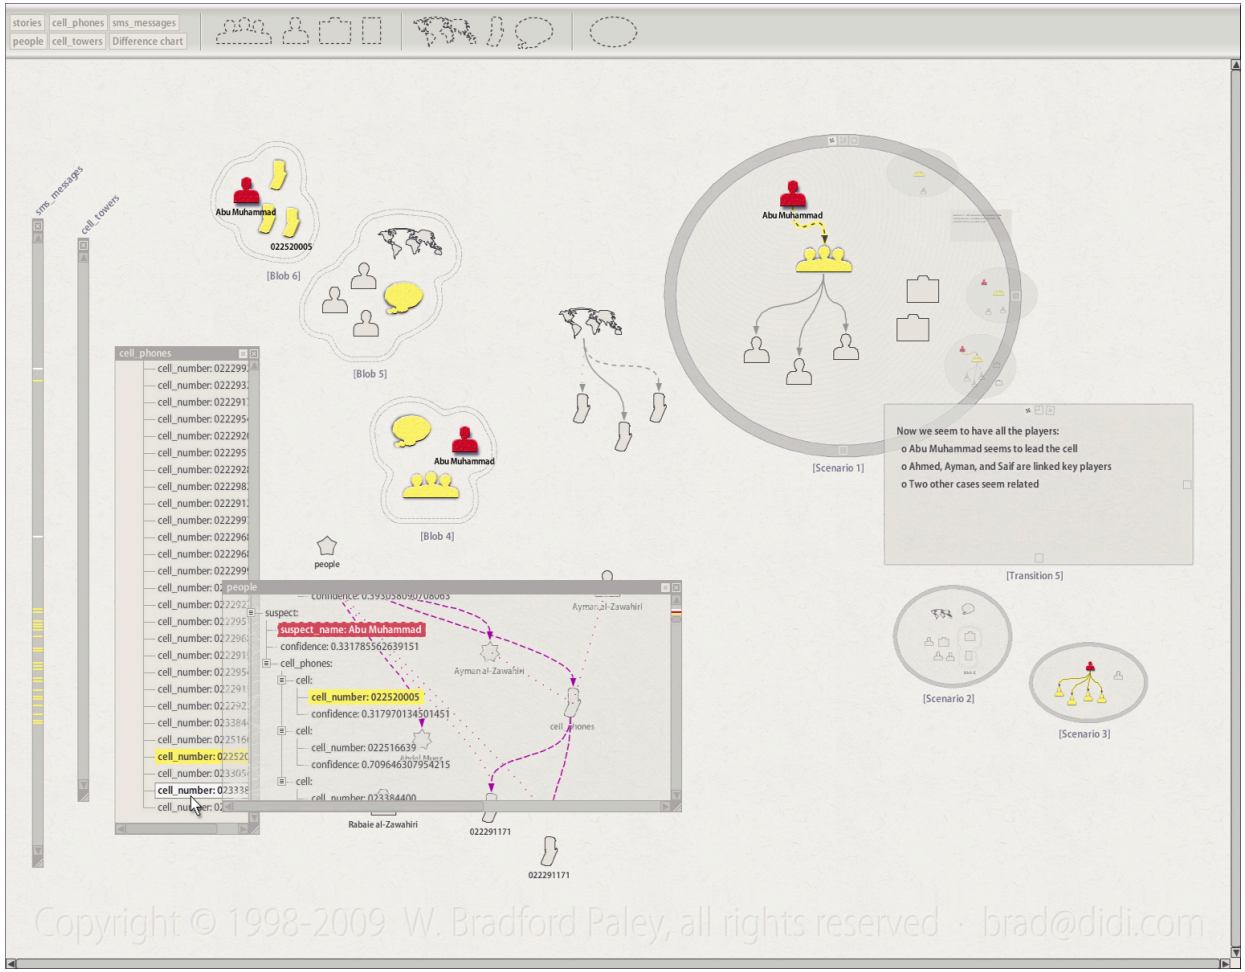

Advanced Charting

Rosenblatt, Penserra Securities, 2009

Key advantages: Integrated interactive Java Applet charts quietly and naturally into an existing Web page; allowed them to be pulled out of the page for closer analysis; innovated point-query and curve comparison techniques based on visual/cognitive principles to reduce misinterpretation risk.

"The Rosenblatt Securities business model requires an innovative, often unique, approach to the trading process utilizing the best technology. Brad not only takes the time to understand our ideas, but creates amazing technology to make them a reality. The result is functional, easy to use, and visually appealing. The fact that Brad is a pleasure to deal with, and works at close to light speed just adds to a positive experience at every level."

—Dick Rosenblatt, Founder

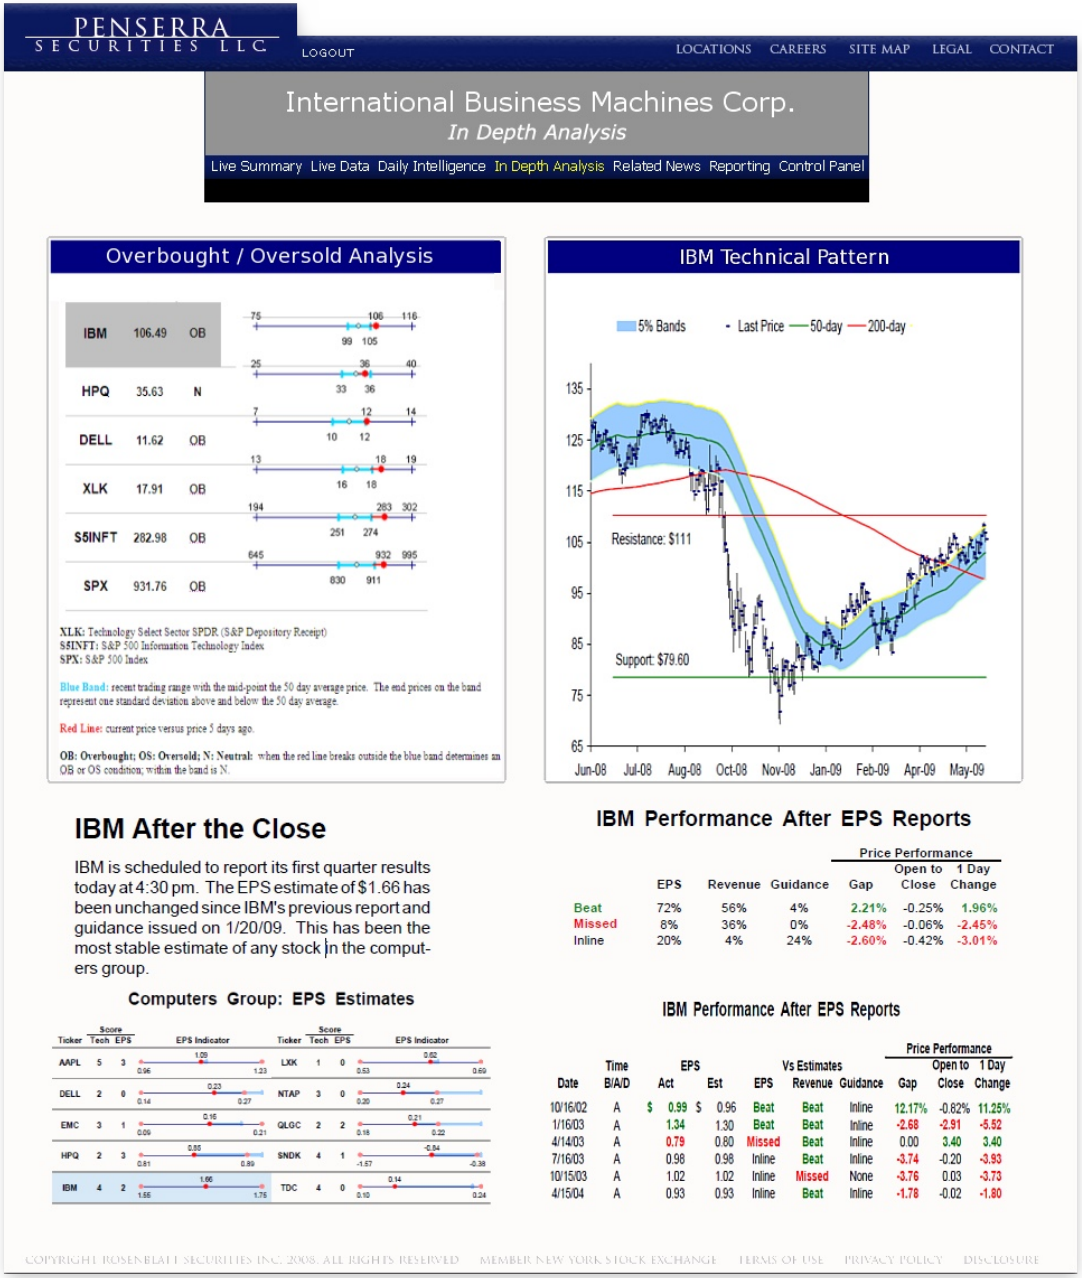

Trade Blotter (Input, Tracking, Alerting)

BIDS Trading, 2009

Key advantage: Our redesigned tools got permanent space on the extremely valuable and generally inaccessible-to-new-tech trader's desktop; before us they needed to be launched for each opportunity—and almost never were.

Three tedious screens/operations were combined into one quick, small tool, which also visually clarified the new transaction-related business process this dark pool innovated.

"We used all the standard conventions broadly recognized at the time, not dissimilar to what's in use today. We were asking a very busy and impatient group to do something new and different. We were not breaking through.

We were at risk of failing.

I met Brad the next year, saw his portfolio, and was astonished. I knew I'd found the key to getting BIDS accepted by our market.

Brad led a redesign effort that resulted in a key client telling us he'd 'never seen an interface improve so much, so fast.'

A few years later, a technology refresh was needed. When an important client was told a new version was coming, he told our team 'By all means—just don't change the way it looks.'"

—Paul Hanson, Founder

"It is a transformational piece of software for us. The first generation was clear and easy, but it was not nearly as simple as it is now… If you were to look at our version 5 and version 6, other than the fact that it says 'Bids' on it, you wouldn't know [it was the same system]."

—Tim Mahoney, CEO

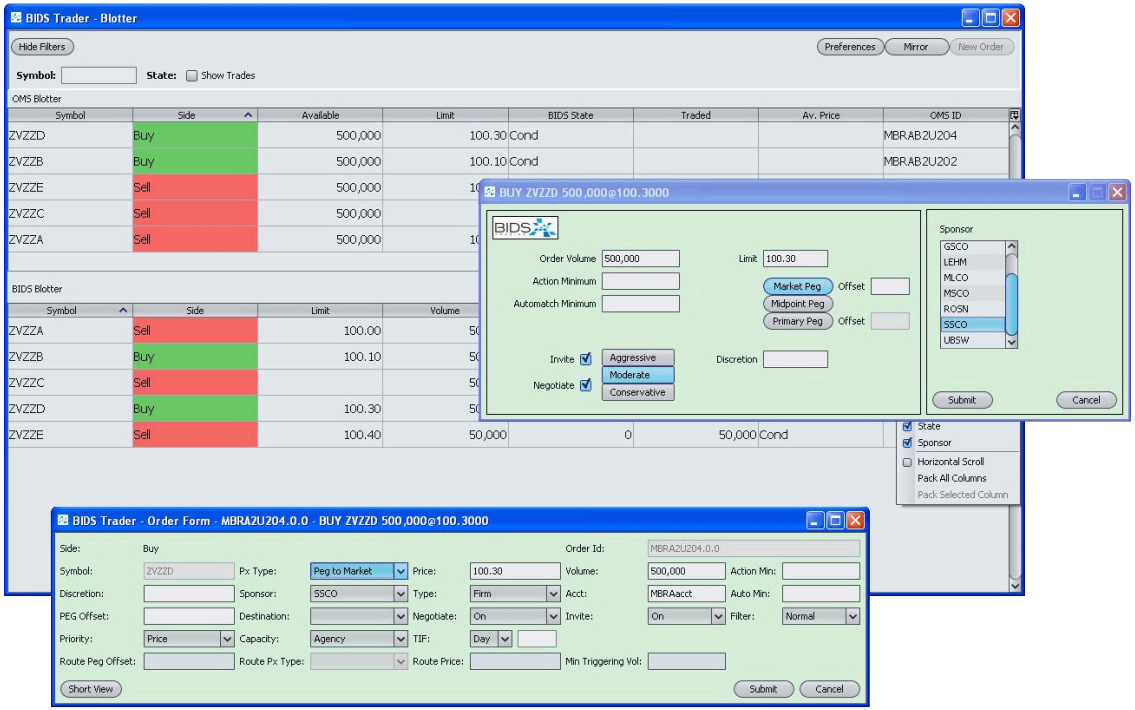

Before

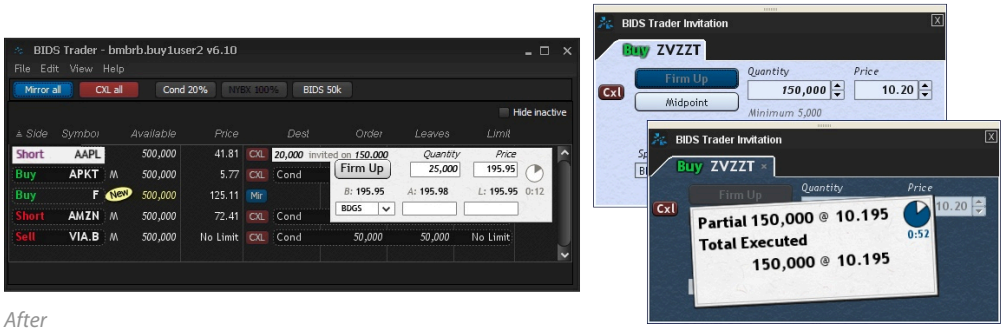

After

FFI (Prototype)

Fixed Income Broker, 2009

FFI is a working prototype of a broker-centric fixed-income trading front end.

Key Update: Improved every individual screen layout to better follow expert trading behavior, adding impossible-to-miss alerting and second-nature multi-key filtering.

Deliverables: A complete system redesign, plus working and tested source code for an enhanced Java widget set that simplified visual presentation, embodying cognitive engineering information layering and other innovations.

XTL Trading Front End

JP Morgan Asset Management, 2010

XTL is a prototype of a complete set of tools for trading equities in "eXtreme" Trade Lists. Every tool has several significant innovations, informed by the traders themselves since we worked for months with them directly on the desk.

The role of traders was changing because of electronic market access and automated trading tools. We anticipated that the number of active/working trades might increase from the then-common thirty or so per trader (which could be handled in a grid: their "blotter") to today's hundreds or thousands (impossible to work in a grid, hence "extreme").

Description: A fully-functional prototype of a game-changing Execution Management System (EMS) front end, expanding to cover Order Management System (OMS) functionality as well: an OEMS for the growing class of "portfolio traders".

Key Cognitive Engineering Change: Equities "consoles," plus other task-specific tools, gather *all the data necessary for a task in one place* and allow all necessary actions to be taken within each tool. This obviates the need for a four- or six-screen setup—most trading could now be done on a tablet. Read the white paper at didi.co about how workflow is hugely optimized by refactoring the trading desktop around the true scarce resource: traders, not today's historical accident of factoring around vendors or data connections.

These tools are direct visual illustrations of the cognitive steps an expert takes to complete a task. As such they look completely different from one another. Pre-attentive recognition allows the next tool needed to be found instantly, and primes in memory the expert's cognitive "scripts" for the next step in their task, facilitating faster and more flow-retaining intellectual task-switching.

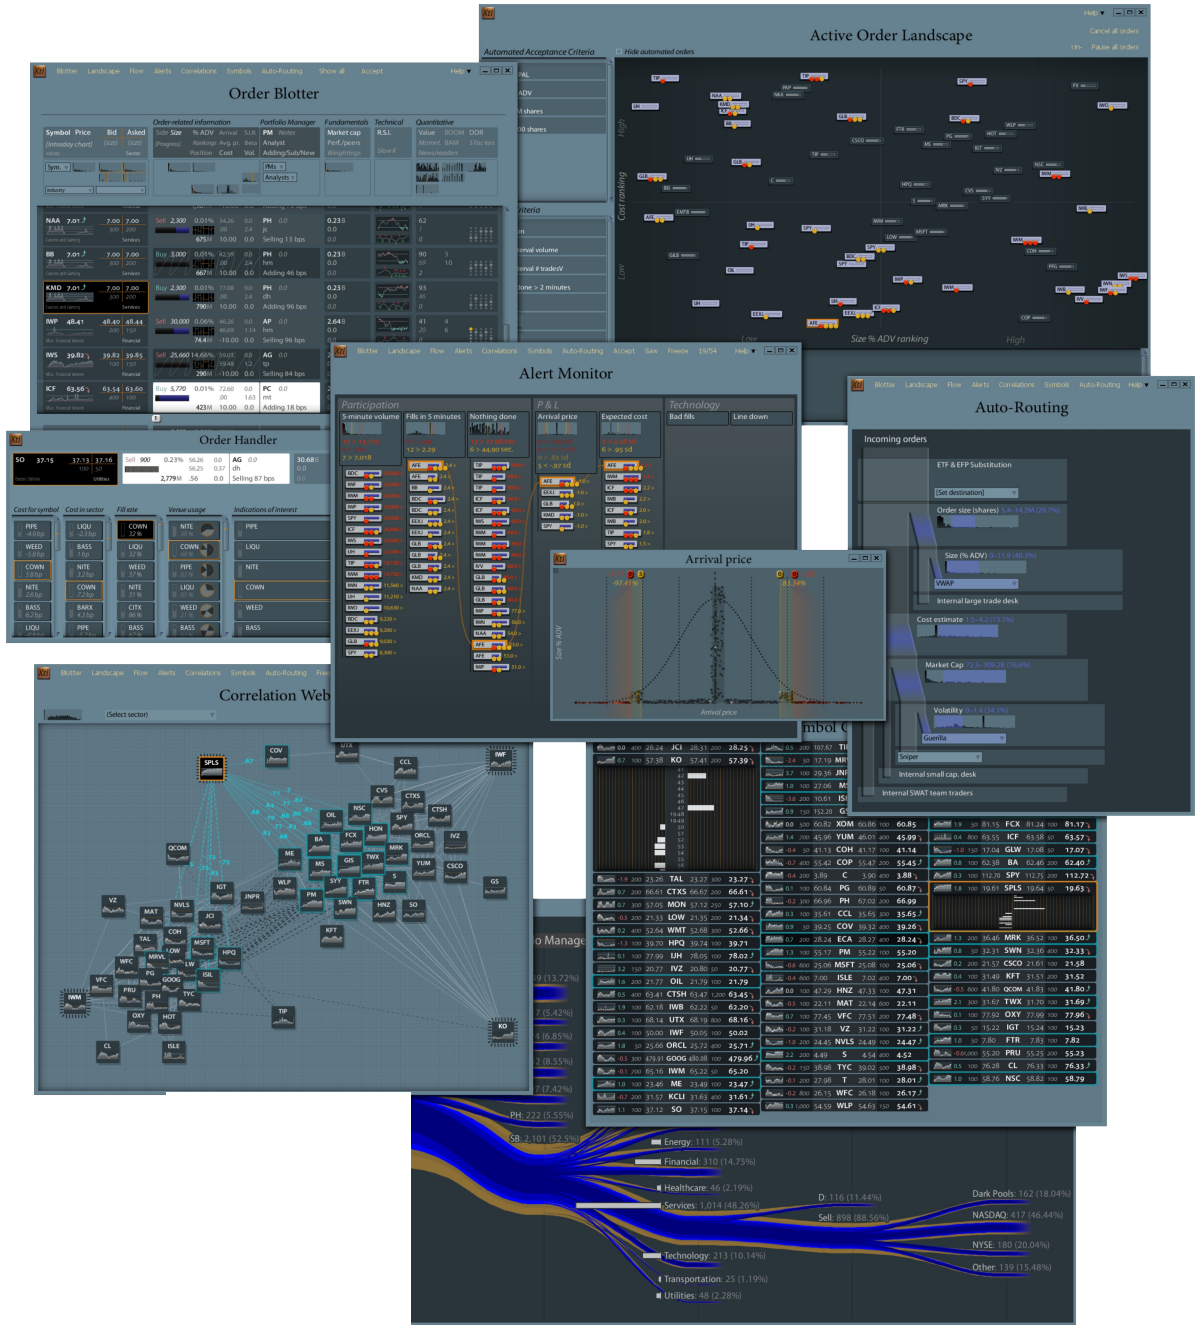

Compass: Execution Management System

Tora Trading, 2011

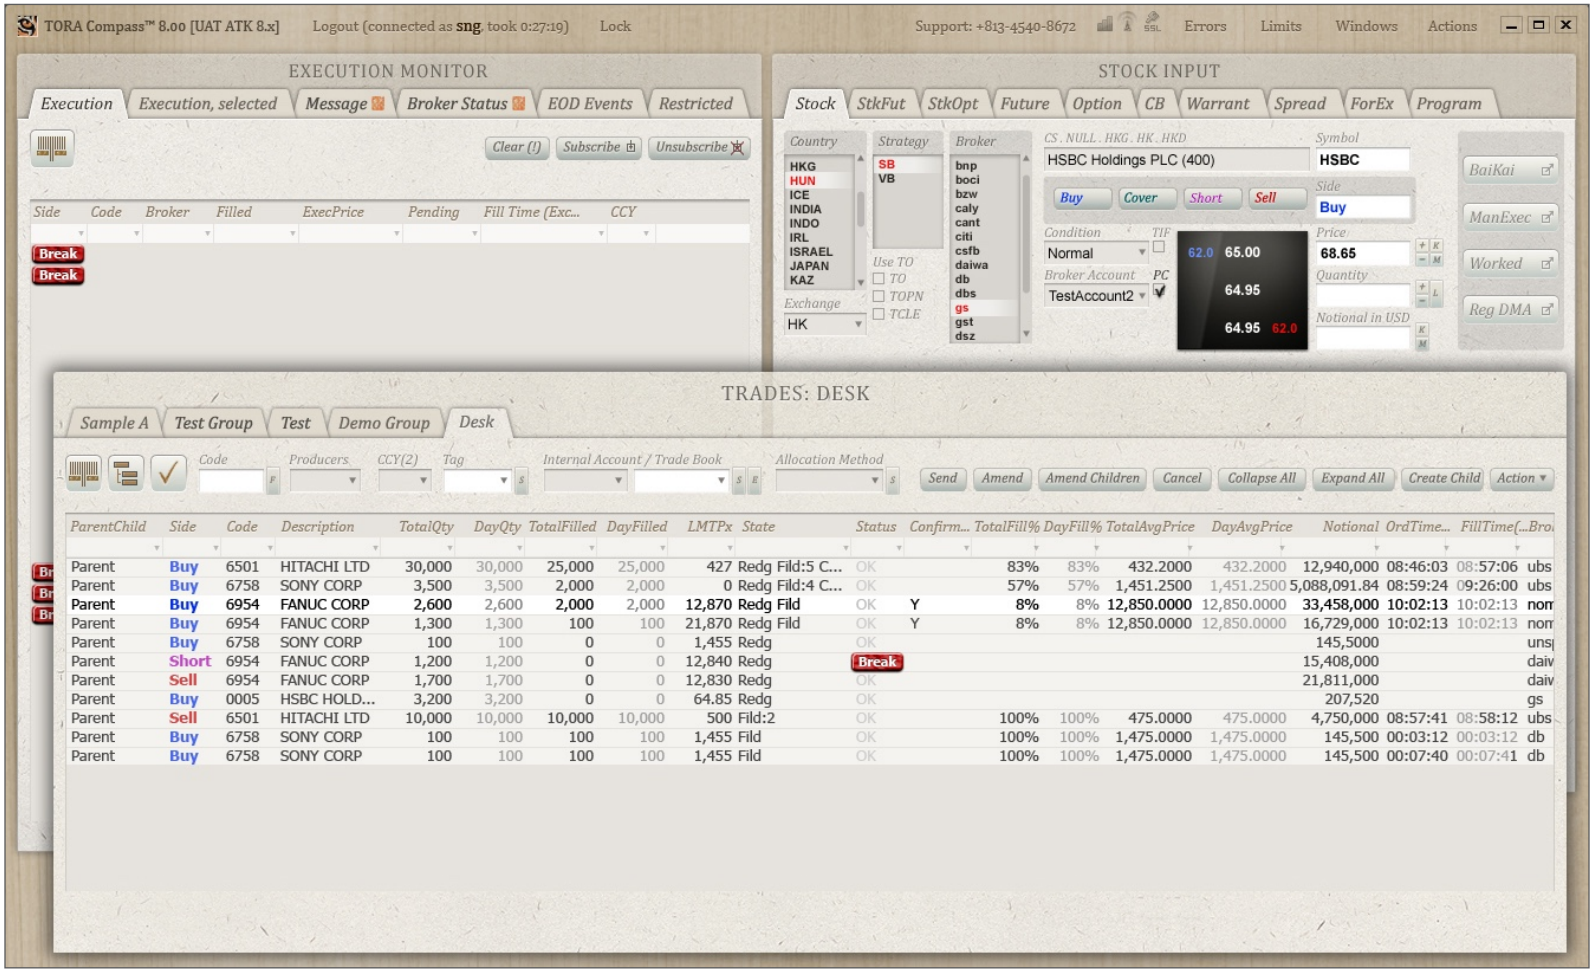

Key advantage: Improved every individual screen layout to better follow expert trading behavior, adding impossible-to-miss alerting and second-nature multi-key filtering.

Deliverables: A complete system redesign plus working, tested code for enhanced Java widgets that simplified visuals, introduced information layering and other innovations.

NetMap, Large-Scale Digraph Analytics

Morgan Stanley and Others, 2011

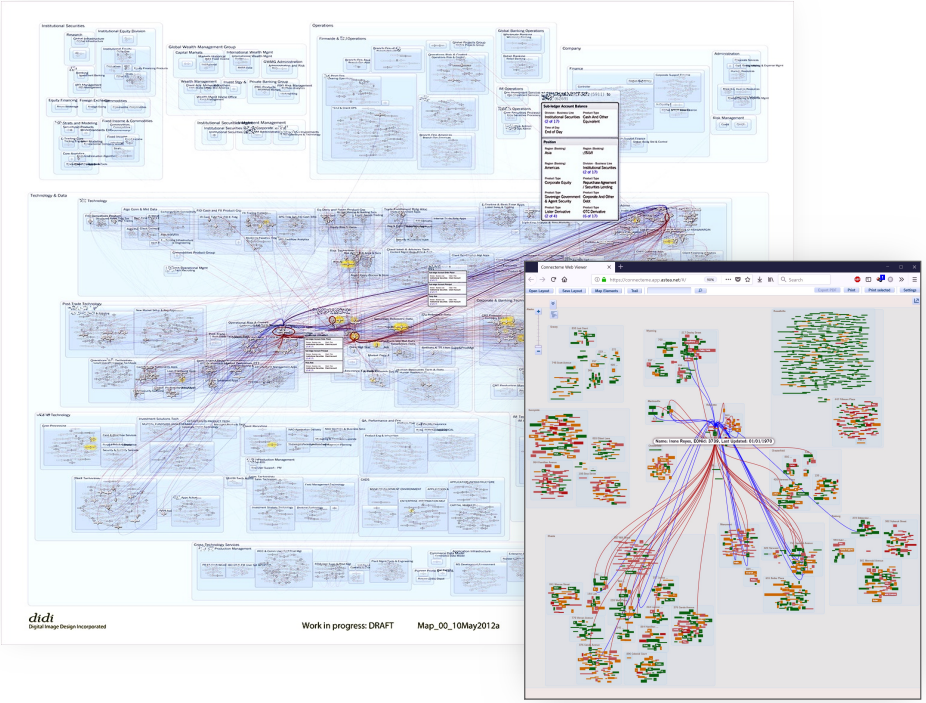

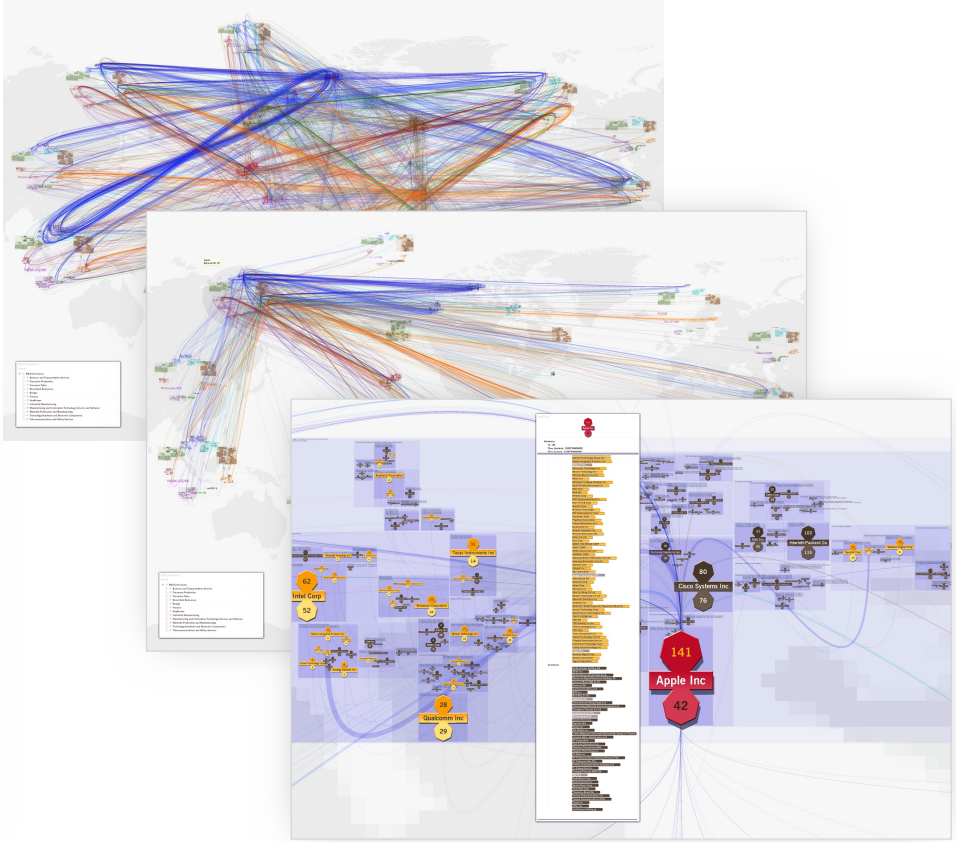

NetMap (previously "Connecteme") displays digraphs: graphs with directed links between nodes. It helps people understand network behavior by showing both overall and detailed structure, inspired by how we're starting to understand the brain's connectome's neural links between synapses.

Above: An early application to data flows between people and between systems, organized into a corporate hierarchy. This proof of concept was so readable it was in front of the Morgan Stanley board of directors within three months of starting.

Right: Companies and their supply chain relationships, organized into a colorful sector hierarchy. Using a world map underlay as a recognizable visual "label" for sector hierarchies in different geopolitical markets, this shows how seeing an overview of some 20,000 links among 6,000 nodes can still support insight and hypothesis generation. A dozen nodes/links is generally unreadable.

One can click on a country to see all the supplier or customer links to and from that country (top: the US, middle: Japan), do the same for any sector or subsector (not illustrated), and zoom in to see an individual company, its sector context and peers, and all of its suppliers and customers (bottom).

Key advantage: Directed-graph networks generally become unreadable hairballs after a dozen nodes and a score of links. Here, thousands of relationships can be interpreted as a gestalt due to hierarchical, human-assisted (therefore *meaningful*) layout of nodes, and one can mentally filter down to make very specific insights about even individual nodes and links, then drill in for specifics.

Key Encoding Innovation: Direction of links in this digraph can be determined at any scale due to the suppliers-in-the-top and customers-out-the-bottom visual encoding.

Key Behavioral Innovation: Nodes are initially laid out with automatic rules, but humans can move nodes or groups to bring conventional wisdom into the layout for better understanding (e.g., put sectors that mainly consume toward the bottom, or a company's C-suite at the top); especially useful when local experts organize their own domains.

Portfolio Management & Risk Analysis

RMS, 2012

RMS was the 100% market-penetration risk modeling platform in the insurance industry; founder Hemant Shah was a primary force in moving insurance underwriting to a quantitative practice, rather than a historic/accumulation model. We were privileged to work directly with him on this and several other projects (like Risk Atlas, next).

Key Transformation: Refactoring of the entire "Cat Analyst" interface, reworking a high-cognitive-load data-centric screen design in favor of task-centric tools. In the process dozens of entities involved in risk analysis were made more concrete and therefore easier to manipulate, think about, and act on. This let screens display and integrate much more information yet almost-paradoxically be more comfortable and evoke more insights.

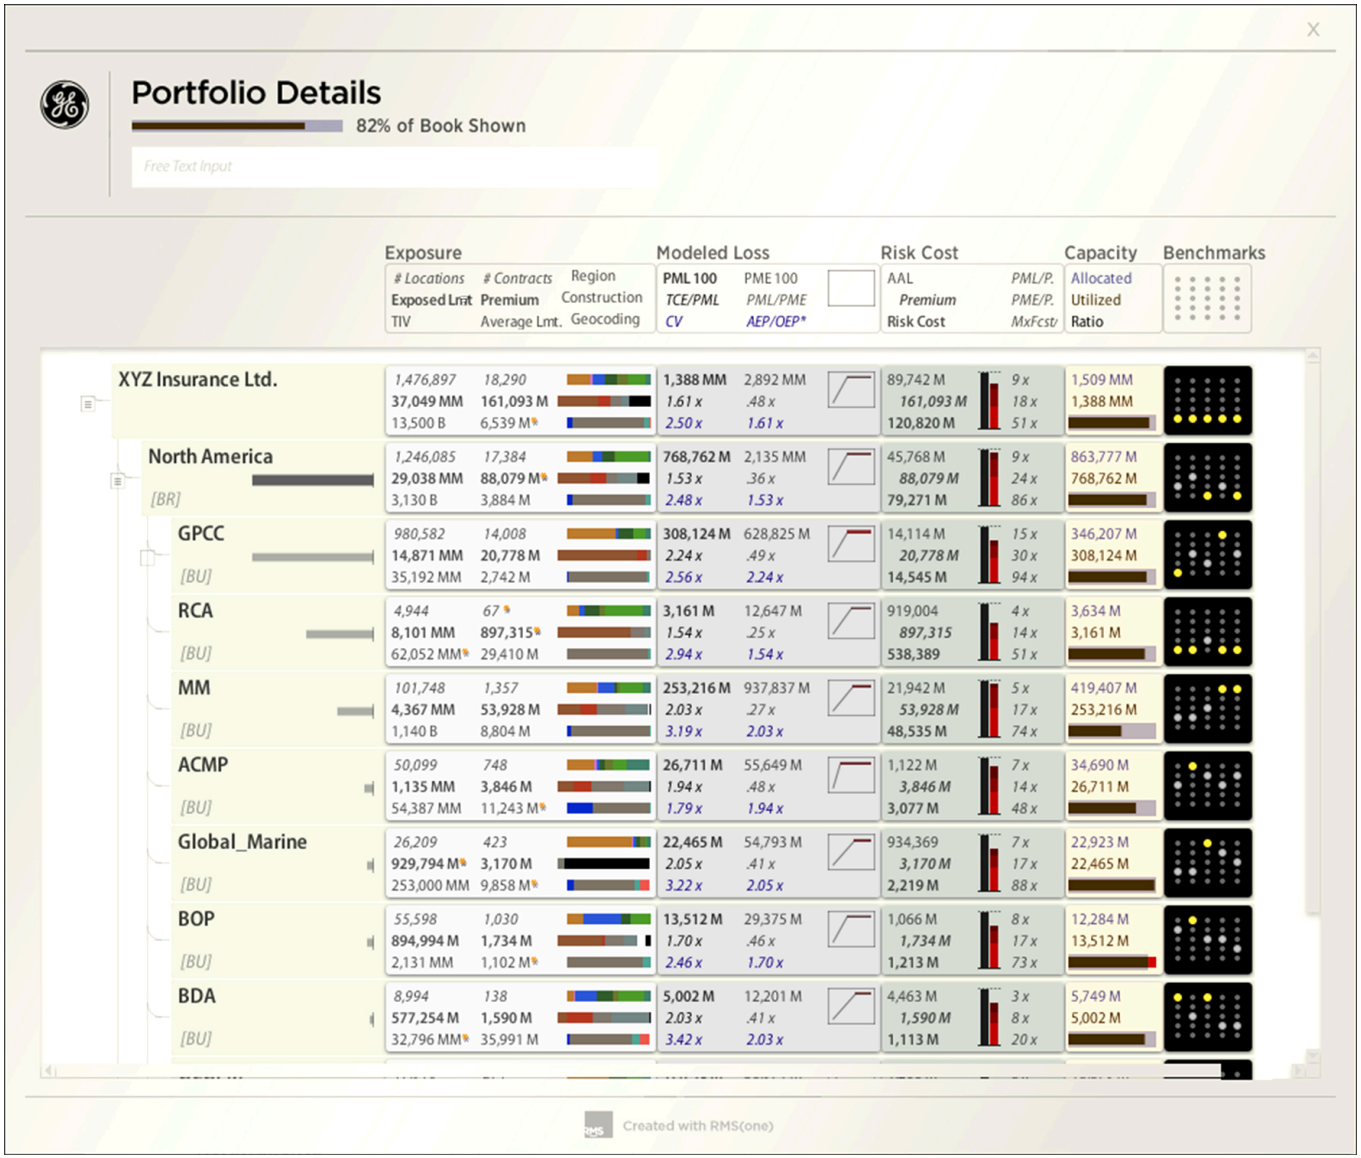

Risk Atlas

DIT2014 (Working LLC Name for a Startup), 2013

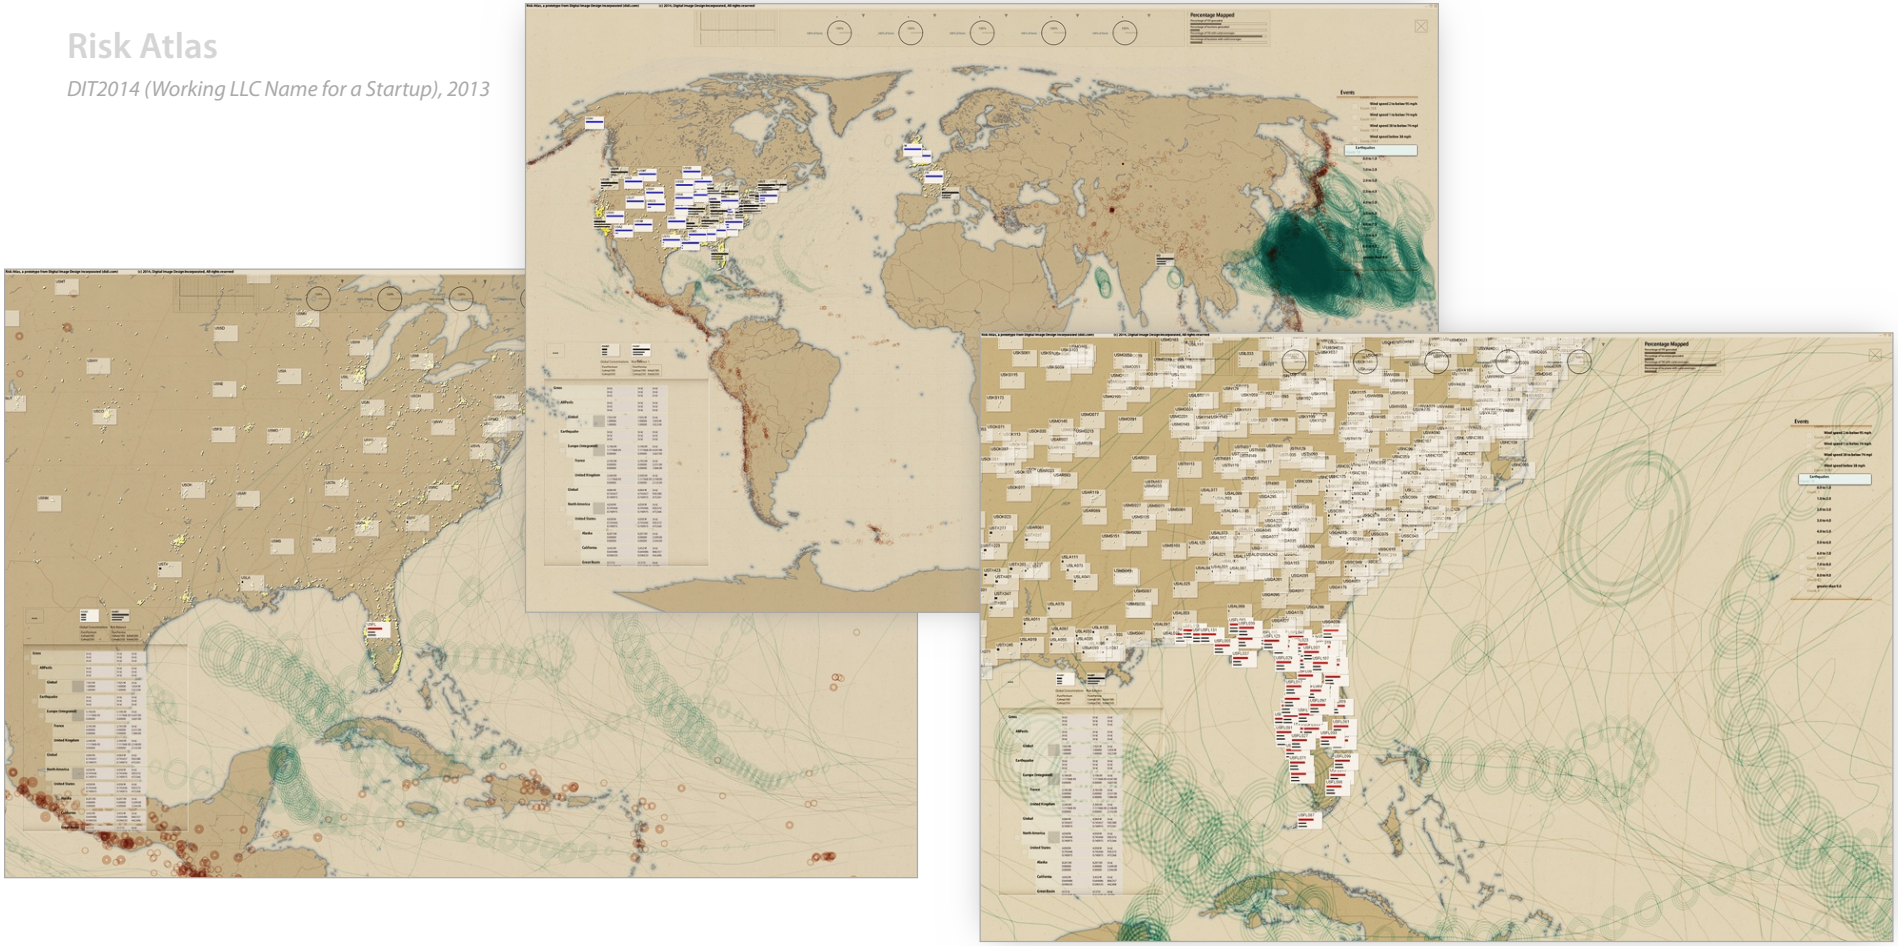

Re/Insurance C-Suite Risk Analytics/Overview

Support for an Overlooked Audience: Allows C-Suite insurance executives to browse billion-element portfolios on a 4k touch screen, sliced and diced by geography, business division, or any other relevant data-cube dimension.

Executives can simultaneously see up to a thousand half-inch glyphs that graphically illustrate strategic business issues, e.g., risk concentration/balance, or premium vs. expenses vs. expected loss: everything needed for each analytic task.

Key advantage: It has never before been possible to break a portfolio into so many cells while retaining the global context that allows comparing details to trends and outliers. The console for each cell can tell a complete business story: custom consoles are added for different tasks or presentation stories.

Key Innovation: Supports the hand-off of questions from executive or board-level management to their direct reports, facilitating faster and more exact communications between colleagues—or people in completely different roles—because all are looking at, and thinking on, exactly the same page.

Energy Trading, Generative Toolkit

Macquarie, 2016

Key Transformation: Rather than slavishly following tidy-making UX design rules, and blowing way beyond the target screen area as the previous designers did, we *listened to how the traders spoke*. This let us give them back the ability to *think like traders*, not be forced fit their thoughts to the screen like clerical data-input staff.

Transformative Interaction Innovation: Our "Composite Fields" are the visual transcription of how experts speak to one another—and therefore *how they think*. Each composite field is a thought "chunk" (a fundamental finding in Psycholinguistics). Each becomes instantly recognizable, so it can be reused and understood in different contexts. The system generated over 30,000 screens on the fly from these intellectual building blocks. Composite Fields are tiny in space and in cognitive load compared to the conventional, wasteful UX approach of our predecessors. And of course, this generalizes: *every* domain of practice has its own recognizable chunks.

Both smaller screens here have all the same information: the previous team's industry-standard approach was systematic, but deceptively (confusingly!) "clean," UX rules applied without thought. Ours is articulated around trading concepts, providing an intellectual scaffolding that helps them execute more quickly, with lower risk—and also develop higher-level insights more easily.

"We've rolled out most of your designs into production, and the feedback has been positive, people booking complex products without having to understand the application been given training."

—Managing Director, unsolicited email 2018

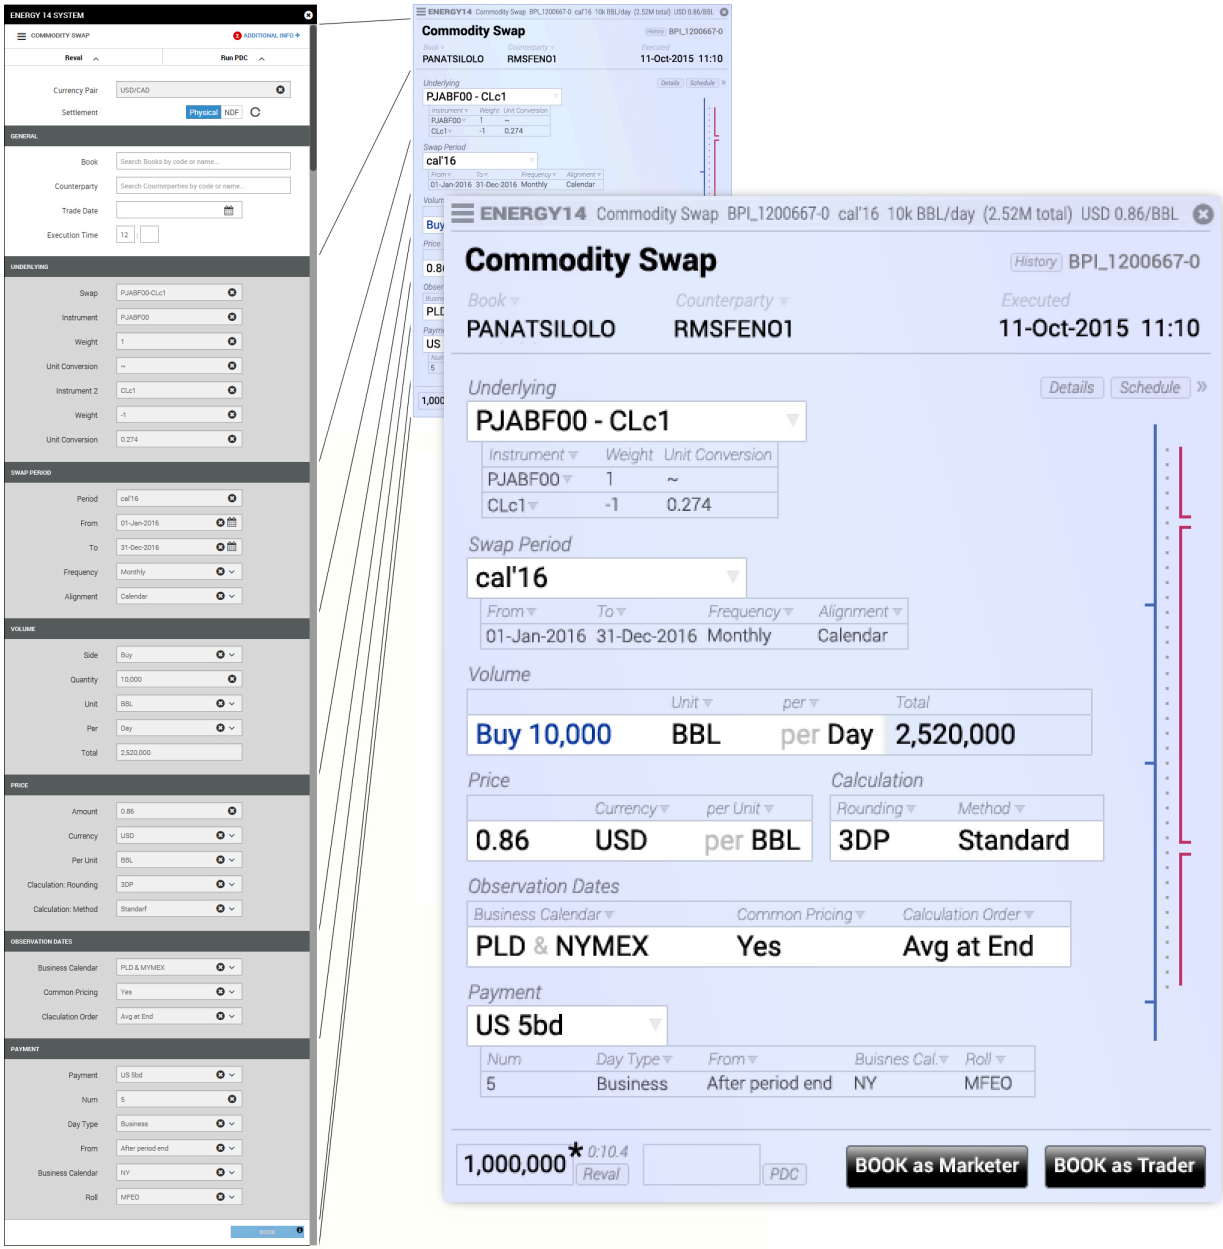

Cogentity

Financial Startup, 2021

Scope: Creation of full company product set, aiming to bring the clarity of Business Mechanics to software startups, then scaling to all corporate classes. We developed all company business strategies & products.

Defining Value: Business Mechanics brings instant, visceral clarity to the operations of an entire corporation, showing relationships between Revenues, Costs, and our newly-introduced concept of Contribution, the essential missing perspective in most startups.

Unique Offering: The Cogentity real-time visual model of an operating corporation allows you, if you're a CEO, to "Discover Your Corporation." It provides concise and accurate overviews of every corporate function, with instantly-available details. These are also available to other roles in the company, helping them cut costs—or bring in more revenues—while contributing to better KPIs, a primary focus for software startups. When shared with investors it can build understanding, trust, and therefore stronger engagement and investment.

Elimination of False Complexity: The video walk-through on didi.co, from a keynote address in 2020, demonstrates how dozens of report pages can be reduced to a single illustration—yet *simplify* understanding. Common presentation tools introduce a layer of dangerous complexity by shoehorning a relatively straightforward story into off-the-shelf presentation modes, like a PowerPoint deck. A corporation is complicated (many moving parts) but not all that complex (the interactions between the parts are known, at least to a first approximation).

An oak tree has *many* more moving parts than most corporations, but it's a pleasure to observe and understand; so are the many moving parts in a corporation when they're illustrated correctly.

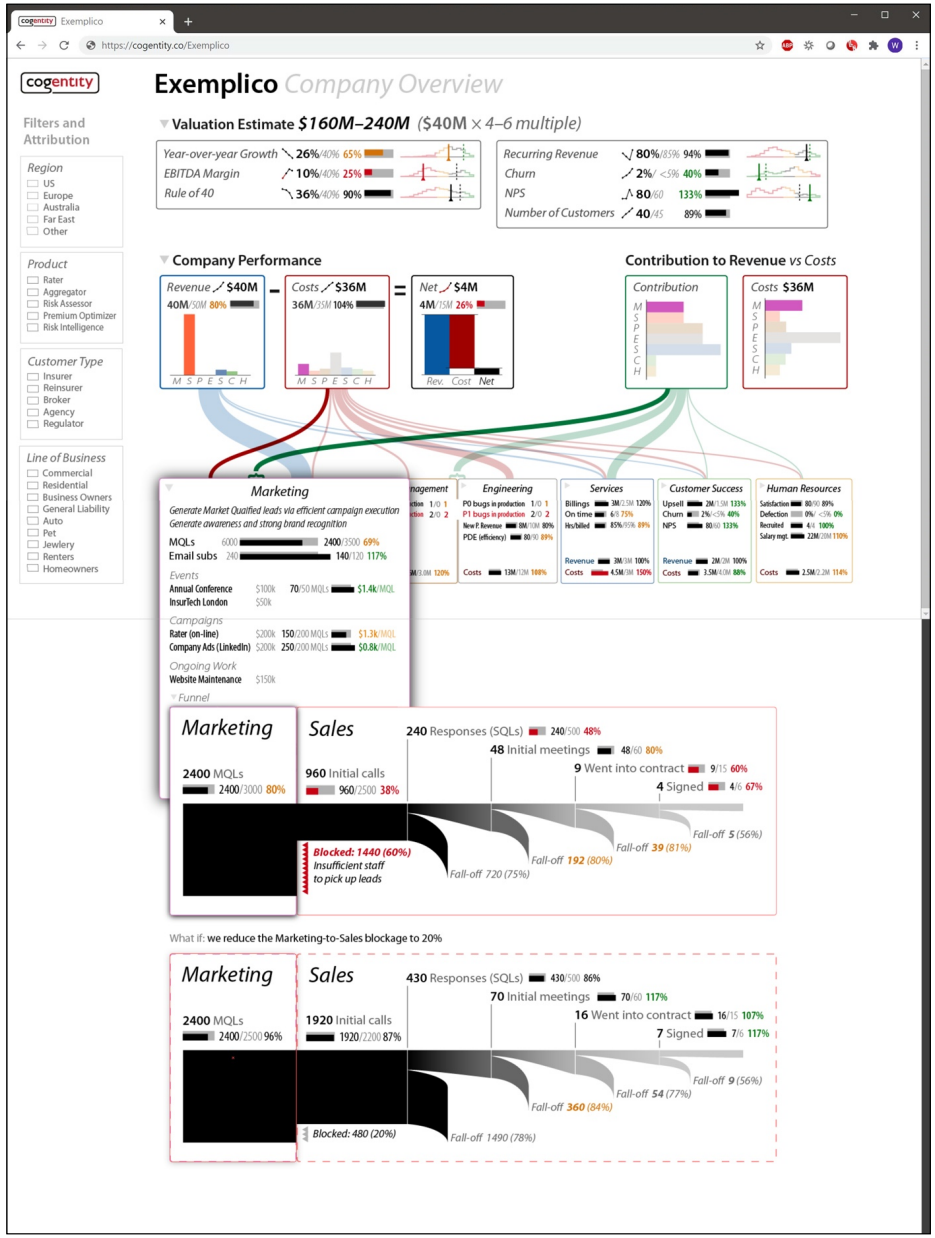

Mosaiq, Unified Investment Portal

Franklin Templeton, 2024

Scope: Complete visual and interaction design for these firm-wide industry-leading portfolio manager and analyst tools.

Process: We designed the entire front end, including all visual context and behavior. Then we did deep dives directly with portfolio managers and quantitative analysts, tying their workflows seamlessly together with tools that embody the understanding and behavior of experts in each role.

Key Advantage: This project is true Cognitive Engineering: an exact fit of expert workflow. But it's played out in a plain-old UX visual context to provide an easy on-ramp for experts used to common Web UX, and to make it easily implementable with the existing tech stack. The critical difference: where plain UX often *squeezes out* domain details—"vanillifies" each domain—to conform to overly-abstract paper-form-like 60s/70s widgets, we *embraced and illustrated* those details to telegraph domain familiarity. These screens *feel like the business*, not the usual lifeless, undifferentiated bunch of widgets experts need to painfully learn by rote—to practice a domain they already know.

Commercials: Thousands of screen mock-ups were delivered over three years; we led the collaborative honing of each one with expert business-side co-designers. To exactly fit the questions and concepts in the minds of experts we injected scores of our own innovative "UX++" affordances. Each was instantly recognizable and usable when deployed, generally with no explicit training at all. The same mock-ups also provided cut-and-paste CSS, speeding implementation and improving code-to-design fidelity. Screens were delivered as 100% client-owned IP, but created from pre-existing didi elements that are available to fit future workflows just as tightly.

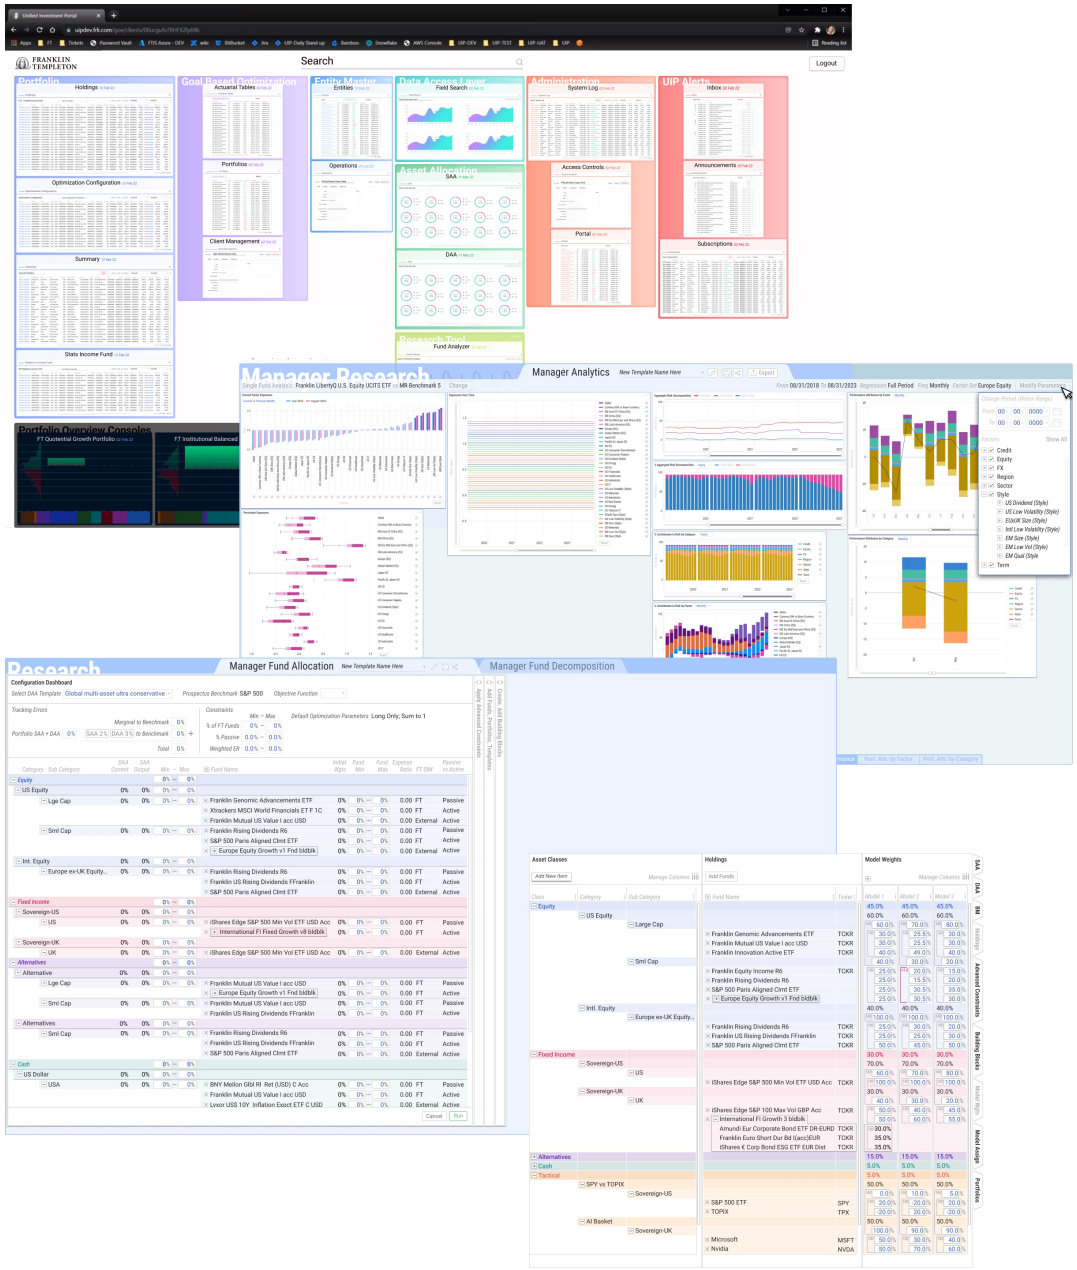

3 Backgrounds

Digital Image Design Incorporated: A micro-CV of Primary Participants

Backgrounds

W. Bradford Paley · Digital Image Design Incorporated Founder, Cognitive Engineer

Analysis/Transaction Front Ends: Goldman Sachs: NYSE broker handheld; Merrill Lynch: Equities EMS; Morgan Stanley: FX EMS for Sales- and Market-facing traders; The New York Stock Exchange: NYSE's own handheld, Specialist Workstation; REDIPlus: Full EMS; BIDS Trading: Dark pool front end; FFI: "TRADEWORKS killer" Fixed Income Broker Tool Suite prototype; JP Morgan Asset Management: Next-gen OEMS prototype; Tora Trading: full EMS front end; Merrill Lynch: Tora's EMS, white-labeled; Macquarie: 30,000+ (!) Energy Trading screens; Franklin Templeton: Complete visual/interaction design framework for Mosaiq, their Portfolio Manager, Analyst, and investor tools; Lehman Brothers: InfoVis for neural networks and other advanced research—in 1989! · *Plus twice as many non-financial systems for similarly global clients*

Academics and Recognition: Adjunct Associate Professor, Columbia University Computer Engineering Dept. (appointed to teach own work as graduate seminar); UC Berkeley Phi Beta Kappa degree in Economics (4.0 in major/minor); ID Magazine Design Distinction Award, Tokyo International Media Festival Grand Prize [Non-Interactive], Ars Electronica production award; NYSCA grantee; NYFA fellow; four US Utility Patents (cited 600+ times so considered a fundamental patent; average citations are ~15)

Publication Features, Selected Book Inclusions: The New York Times (articles on TextArc, CODeDOC), Nature, Geo, Slashdot, Esoteric.codes, Visual Complexity (Website & book), Rattle (the tools as visual poetry), ID Magazine (Design Distinction award, feature on the NYSE handheld), Atlas of Science, Database Aesthetics: Art in the Age of Information Overflow, Living With Art, Digital Art, Wired, PBS Science Friday in-depth interview

Museums and Galleries: The Museum of Modern Art, The Whitney Museum of American Art (permanent collection), ASU Gallery (solo show), The Chelsea Art Museum, Streaming Museum, Pace Digital Gallery (Pace university), Cooper Hewitt, Smithsonian Design Museum (permanent collection, upcoming), People's Republic of China Central Academy of Fine Arts Museum

Keynote/Plenary Addresses: Graph Drawing, Diagrams, 9th International Conference on Information Visualization ("InfoVis"), SmartGraphics, 2020 Summit for Astea/FMI ("Bulgaria's MIT"), Institutional Investor's Financial Technology Forum, UCLA IPAM's Social Networking, IEEE InfoVis, ALT.CHI, The Kitchen · *Plus scores of invited talks, presented to thousands of scientists, and workshopped in seminars and classes with hundreds of students*

Warren Kemp · Creative Director, Product Designer

Analysis/Transaction Front Ends: Davidson Kempner: Deals Snapshot, Morgan Stanley: FX EMS, NYSE: Specialist Workstation, Tora Trading: EMS, Franklin Templeton: Mosaiq

Select Studios, Agencies: Atmosphere BBDO, Ixonos, Ogilvy & Mather, Publicis Modem, Rapp Collins, Tractor Studios, Tribal DDB, Wunderman

Select Clients: Al Jazeera America, Cingular Wireless, Citi, Consumer Reports, Emirates Airlines, Fair Food Program, GE, IBM Brand, IBM Printing Systems, LG Electronics & Appliances, New Yorker Magazine, Pacific Road Capital, Starwood Hotels, TCL Electronics, The Years Project, Zuckerberg San Francisco General Hospital

Awards and Recognition: Andy Awards, Cannes Lions Advertising Awards, The Art Directors Club, The Clio Awards, D&AD Awards, The London International Advertising Awards, The One Club, The Webby Awards, World Luxury Award - Monaco

Peter Kennard · Infrastructure Architect, Cognitive Engineer

Interactive Product Design, Finance: Citibank 1981-2, MarketMate: eBay-like venue for financial instruments and general merchandise; Mansfield Charts 1979-80, perhaps the first rigorous on-line chart/data service for ~6000 mostly technical-analysis brokers and traders

Interactive Product Design, Animation: Calgari TrueSpace 1986-7, partner in the first true 3D interactive modeler: architected, coded the kernel & interactive front-end, managed dev. team, ported to Windows and expanded in 1990-92; Autodesk Animator Pro interactive animation tool 1988-9, primary architect, designer & kernel coder, dev. team lead

Interaction Architect & Development, Fine Art: Dialog Table 2002 deployed at the Walker Art Center; TraceEncounters: a 2.25" wearable pin that discreetly logged interactions among all 1000+ attendees at the 2014 Arc Electronica and drew a real-time social graph (developed and tested a tiny real-time interrupt-driven multitasking OS on the pin's CPU in two weeks)

Miscellaneous: Smithsonian Museum 1973-4, printed from Ansel Adams negatives; Bergdorf Goodman 1977-9, lighting for their renowned window displays; member of Andy Warhol's entourage, visited the Factory