THE PENTIUM CHRONICLES

THE PEOPLE, PASSION, AND POLITICS BEHIND INTEL'S LANDMARK CHIPS

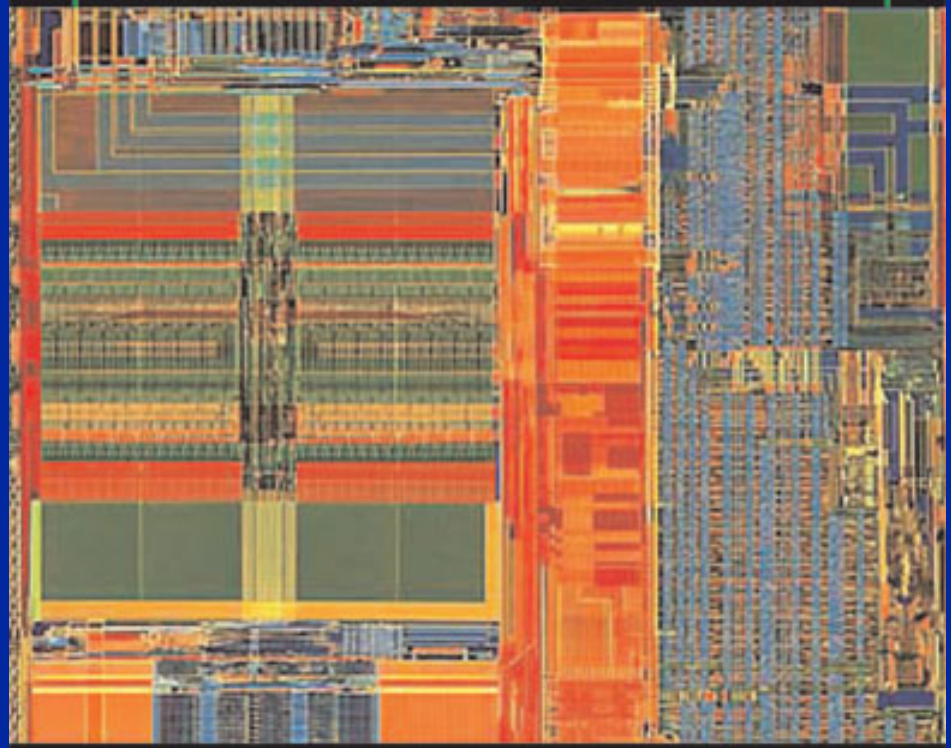

ROBERT P. COLWELL

Aims & Disclaimers

1. Engineering progress comes from real-world learning
2. Following is my opinion only. Not sanctioned, approved or checked by Intel.
3. I am not unbiased. I'm proud of what Intel Corp. has accomplished, and proud of our team's contributions to it.

Engineering Basics

- **Scientific method**
  - Most brilliant innovation in human history
  - Hypothesis, experiment, conclusion
  - Repeatability & control are keys
- **Engineers don’t have that luxury**
  - Best we can do is learn from failures & successes

Engineering & the Scientific Method

- Engineers strive for data-driven decisions
- But big decisions cannot *be* data driven
  - Kelly Johnson: “I can make a spy plane go Mach 3.5+”
  - NASA committing to moon landing in <10 years
  - Intel commissions team to do new microarch ($B bet)
- Problem is much harder today
  - Economics of silicon require huge volumes
  - Huge volumes means buyers are average public
  - Average public means high-tech subject to fashion whims

    Ring tones? Instant messages? Who *knew*?

No Controls on Engrg "Experiments"

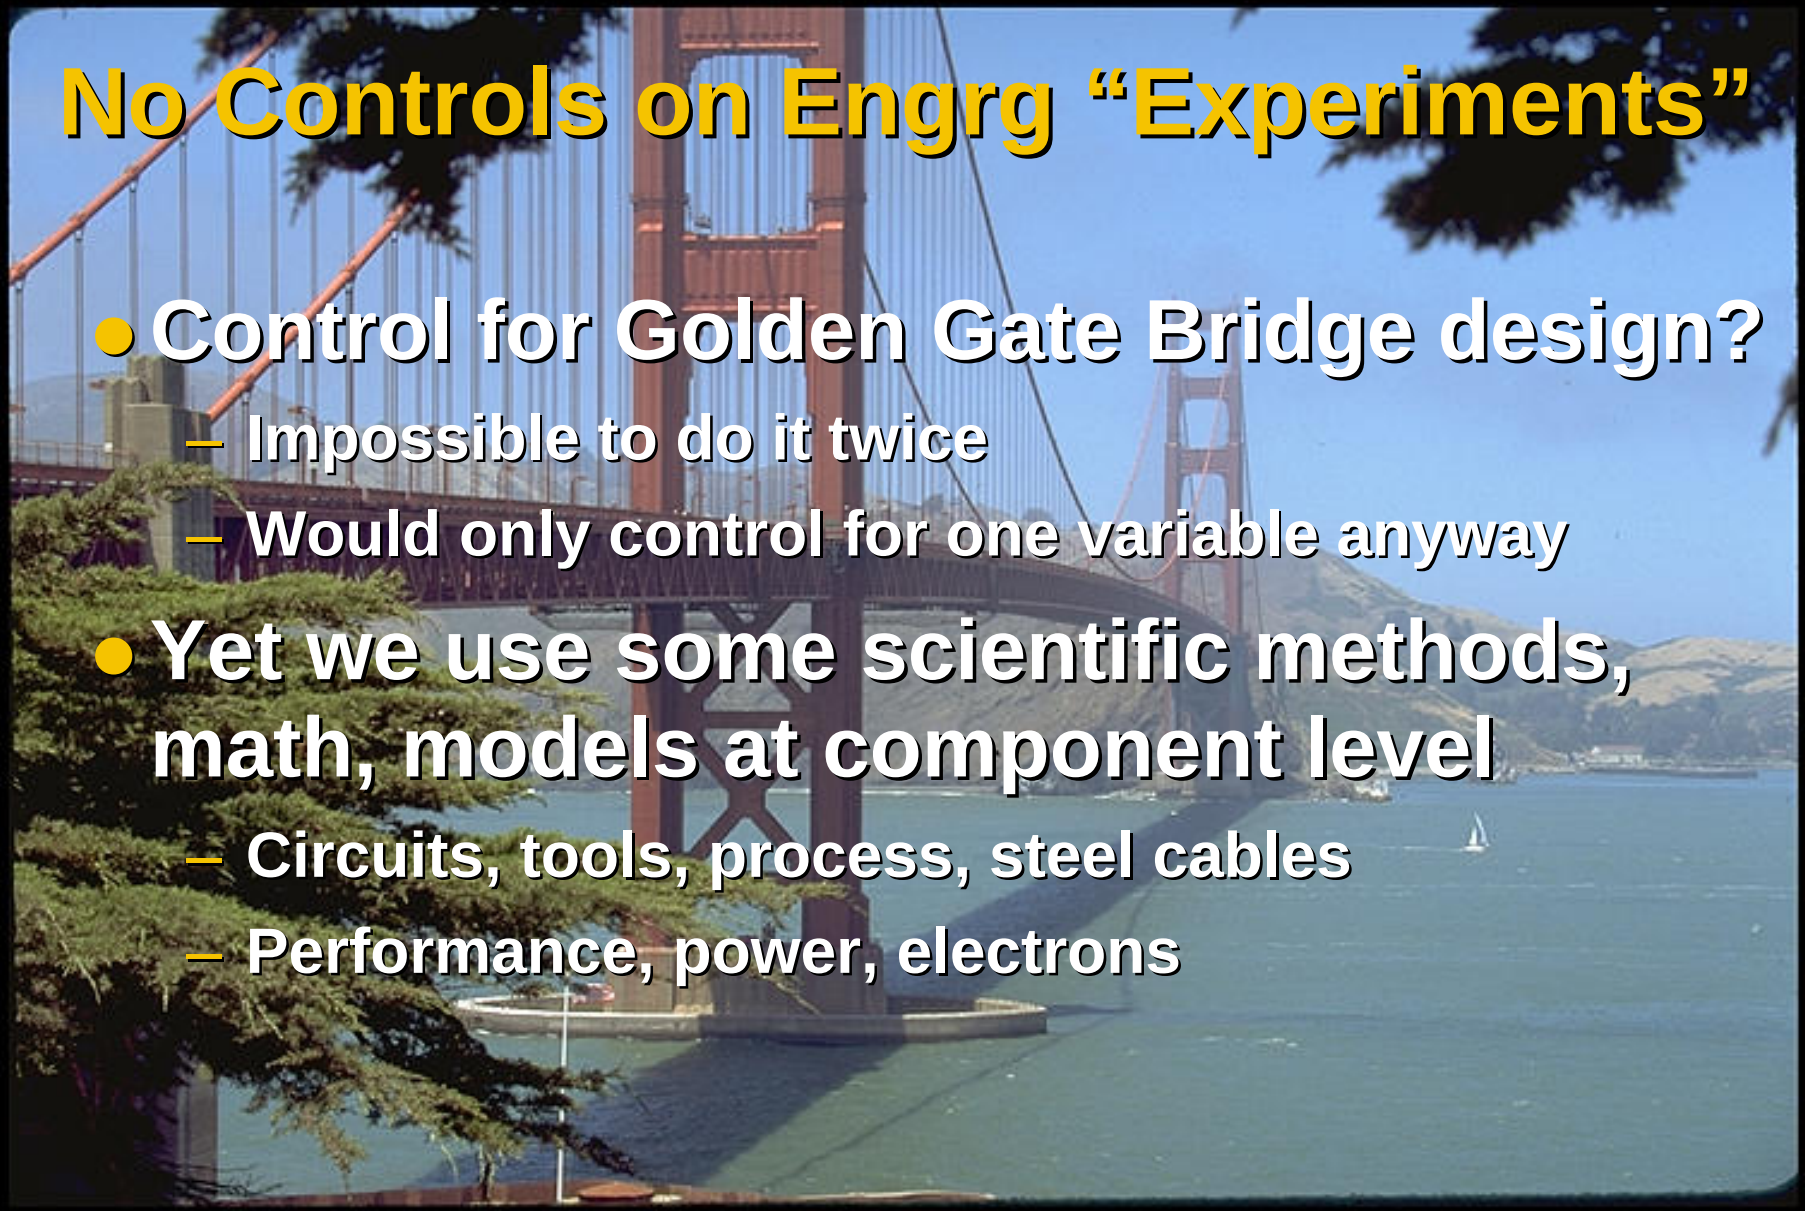

- Control for Golden Gate Bridge design?
  - Impossible to do it twice
  - Would only control for one variable anyway
- Yet we use some scientific methods, math, models at component level
  - Circuits, tools, process, steel cables
  - Performance, power, electrons

Engineering is not Science*

- Must pick right primary targets
  - Perf, features, schedule, power, die size
- Watch for impact on secondary (implied) targets
  - Product quality, longevity, new markets
- Engineers are human
  - Teams, pecking orders, career aspirations, likes/dislikes

* and vice versa

Outline of this talk

- P6 origins
- Project Phases
  - Concept
  - Refinement
  - Realization
  - Production
- People & Intel

Concept Phase

- Setting a data-driven culture
  - Koosh-balls to discourage pontificating
  - Meeting in closet
    - No interruptions, no clean whiteboards
  - DFA tool
- All design teams are not equal
  - Workable concepts must align w/ team
- Collecting customer inputs
- Wizards come in all sizes

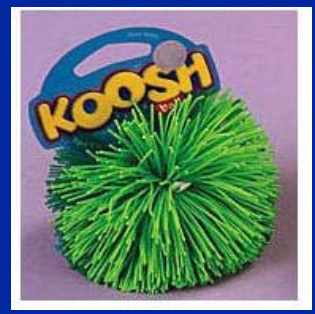

Concept: Origins of P6

- **What was world like in 1990**
  - **RISC/CISC wars & religious visions**
    - **Slow floating point**
    - **Common belief that fastest ISA would win**
    - **Claims of x86 declining**
      - Accumulating old instrs
      - Poor quality compilers
  - **Alliances & consortia**
  - **No internet, only standalone desktop apps**
  - **Buses, platforms**
    - **Reuse Pentium bus?**
    - **L2 cache, backside bus…Compaq**

Concept: origins cont…

P6 architects believed…

- Compatibility paramount
- x86 not necessarily barrier to high speed
  (RISC folks certainly didn't agree)
- Must plan for proliferations
- 32-bit for 10-15 yrs, then 64-bit everywhere
- Could design for servers and waterfall down

Concept: how to start?

How do you conceive new uArchs that take maximal advantage of new Moore's Law transistors but don't overrun that budget or drive you nuts with complexity, yet balance clock speed properly?

Would out-of-order work?

- **IBM 360/91 not reassuring; ran code but not clear how fast**
- **We liked promise of speed w/ object code compatibility**
  - **VLIW experience: fast, simple machines at cost of serious compilers, recompilation**
- **Frank Smith's microdataflow idea got us started, Andy Glew's enthusiasm helped**

Concept: OOO

- Easy to design machines you can never get right*
  - Projected 10M transistors; actually got 5M
    - Ameliorated by L2 cache in package
- Non-constraint: power/thermals
  - Even though Pentium 13W was eventually an issue for Sequent…
  - P6's 23W wasn't a big challenge
    - We just made chip monitor its temp as a fail-safe

* John Hennessy, ISCA '96

Concept: Compatibility

What's it mean to be x86 compatible?

- Intel had no functional simulators in 1990…so Dion Rodgers wrote one
  - Can't do reasonable pre-silicon validation without it
  - Unclear how to make rational implementation trade-offs if you don't know when you've broken something
  - Drives out misunderstandings & fuzzy thinking

Refinement Phase

- Task of Refinement phase is to narrow down 2 or 3 Concepts
  - Must begin to focus on fewer ideas
  - Intent of Ref't phase is to settle on one
- Gotchas
  - Handling non-quantifiables in a data-driven culture
    - Decisions = f(data, intuition)
    - Include risk & take seriously even though non-quant.

Refinement: models

- Behavioral models
  - Expose good questions, provide some answers
  - Reinforce data-driven culture
  - Develop best ideas when still time to backtrack if necessary
    - Good thing we never needed to w/P6; BRTL turned out to be heavy-weight, not quick & dirty as hoped
    - Nowadays you'd try a Matlab model to get quick read on what looks promising

Refinement: ECO's

Engineering Change Orders

- When to put project under ECO control
- Role of czar
- Disallow pocket vetoes
- Should have POR under ECO control but…
  - Didn't try on P6
  - Didn't succeed on Pentium 4
    - Few give up power willingly (real or perceived)
  - Upshot was too many changes per unit time
    - Too fast for Nyquist sampling rate of design team

Refinement: Focus Groups

Bridging architecture and design

- **A gap here could be fatal later**
  - **No throwing over wall**
  - **Focus groups**
    - solved Arch->Design communication need
    - also solved my lousy person-hour estimate on BRTL
  - **Real world focus keeps architects honest**
    - No writing unreasonable checks that designers must cash
- **P6 superpipelining from chance discussion with ALU engr**

Refine: Targets & Techniques

Setting product quality targets

- "From the beginning: perfection" was official motto of a chip designed elsewhere in 1980
- Took 13 respins to get it functional
- This is not the place to "ask for C" if you really need B

Avoid/Find/Survive bug model

- Avoid bugs by good design practices
- Thorough validation plan: find errors that get in anyway
- Survive bugs that escape into wild

Design reviews

- How to, how not to
- My chipset design review gaffe

Realization Phase

Perfect A, mediocre B,C,D

- Engineers must seek balance
- Reaching for “brass ring” on one target out of many is recipe for disaster
- Seek elegance & grace in design, not home runs

Gratuitous innovation

- Engineers are measured by profits from designs, not #patent plaques on walls
- Known-good reuse beats new-and-untried

  Innovate only where necessary

Realization cont.

- **Awards**
  - Very strong motivator, esp. peer recognition
  - Must not reward “firefighting by arsonists”
- **The Simplification effort**
  - Complexity sneaks up on you a little each day
    - Like ice on wings of airplane
    - A little doesn’t matter; a lot crashes plane
  - Take project timeout midway, put search for simplicity at top of project priorities
    - Can’t do this continuously, but reminder helps

Production Phase

- **Functional correctness @ rated perf**
  - **You wouldn't test a race engine by idling it**
    - **Not enough to show functionally correct**
    - **Not enough to show perf goals reached**
    - **Must do both *at same time***
    - **Nowadays must also modulate thermals**
  - **Prove design delivers at worst corners of design space**
- **Find and handle surprises**
  - **Complexity breeds fragility**
  - **Fragility breeds surprises**
  - **Surprises are bad**

People Factor

- **Burnout**
  - 4+ year projects with 6-month crunch at end are tiring
  - But stress & frustration are what cause burnout
  - Successful projects relieve both
- **Treat engineers well**
  - Don't make example of errors or you'll get no more risk-taking
    - Intel did this right with FDIV
    - Do you know the names of those involved? Me neither. That is a *good* thing for all concerned.

Intel-isms

- The Led Zeppelin incident
- Exiting the bag check
- Team building exercises (sailing, paintball)
- Andy Grove's morale booster

Final Thoughts

- This is what we did
  - Not arguing this is best
  - But much of it was pretty darned good
  - P6 was fastest chip in world for several months
  - Most successful general uArch ever
- Have courage of your convictions
  - But flexible enough to change your mind
- Do whatever it takes to succeed

Questions?