CS514: Intermediate Course in Operating Systems

Professor Ken Birman
Krzys Ostrowski: TA

Perspectives on Computing Systems and Networks

- CS314: Hardware and architecture
- CS414: Operating Systems
- CS513: Security for operating systems and apps
- CS514: Emphasis on "middleware": networks, distributed computing, technologies for building reliable applications over the middleware
- CS519: Networks, aimed at builders and users
- CS614: A survey of current research frontiers in the operating systems and middleware space
- CS619: A reading course on research in networks

Styles of Course

- CS514 tries to be practical in emphasis:
  - We look at the tools used in real products and real systems
  - The focus is on technology one could build / buy
  - But not specific products
- Our emphasis:
  - What's out there?
  - How does it work?
  - What are its limits?
  - Can we find ways to hack around those limits?

Our topic

- Computing systems are growing
  - ... larger,
  - ... and more complex,
  - ... and we are hoping to use them in a more and more "unattended" manner
- Peek under the covers of the toughest, most powerful systems that exist
  - Then ask: What can existing platforms do?
  - Can we do better?

Some "factoids"

- Companies like Amazon, Google, eBay are running data centers with tens of thousands of machines
  - Credit card companies, banks, brokerages, insurance companies close behind
  - Rate of growth is staggering
- Meanwhile, a new rollout of wireless sensor networks is poised to take off

How are big systems structured?

- Typically a "data center" of web servers
  - Some human-generated traffic
  - Some automatic traffic from WS clients
- The front-end servers are connected to a pool of clustered back-end application "services"
- All of this load-balanced, multi-ported
- Extensive use of caching for improved performance and scalability
- Publish-subscribe very popular

A glimpse inside eStuff.com

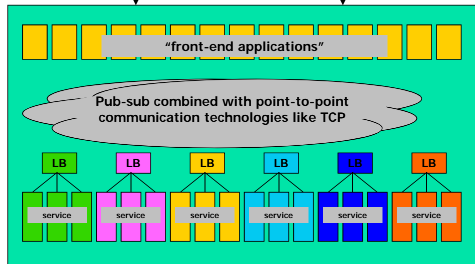

Industry trend: Web services

- Service oriented architectures are becoming the dominant standard in this area
- But how well do the major platforms support creation of services for use in such settings?

Let's drill down...

- Suppose one wanted to build an application that
  - Has some sort of "dynamic" state (receives updates)
  - Load-balances queries
  - Is fault-tolerant
- How would we do this?

Today's prevailing solution

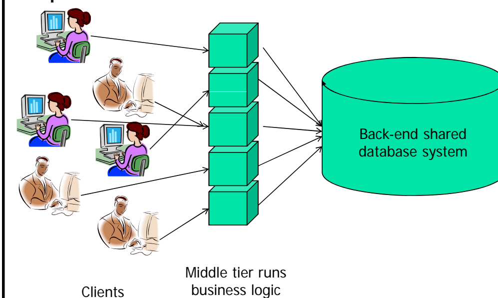

Concerns?

- Potentially slow (especially during failures)
- Doesn't work well for applications that don't split cleanly between "persistent" state (that can be stored in the database) and "business logic" (which has no persistent state)

Can we do better?

- What about some form of in-memory database
  - Could be a true database
  - Or it could be any other form of storage "local" to the business logic tier
- This eliminates the back-end database
- But how can we build such a thing?

A RAPS of RACS (Jim Gray)

- RAPS: A reliable array of partitioned subservices
- RACS: A reliable array of cloned server processes

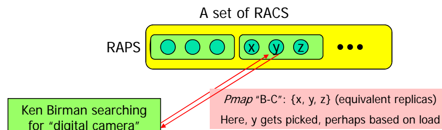

RAPS of RACS in Data Centers

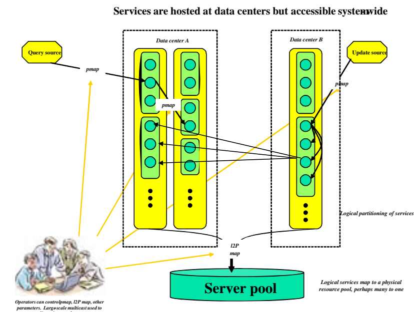

Scalability makes this hard!

- Membership
  - Within RACS
  - Of the service
  - Services in data centers
- Communication
  - Point-to-point
  - Multicast
- Long-distance links
- Resource management
  - Pool of machines
  - Set of services
  - Subdivision into RACS
- Fault-tolerance
- Consistency and monitoring mechanisms

Technology needs?

- Tools to build scalable services lacking today!
- Web services
  - Standardizes the client – data center path
  - But *treats the internal structure of the data center as a black box*
- Three-tier middleware (databases) can help
  - But some applications don't fit this model

More technical barriers

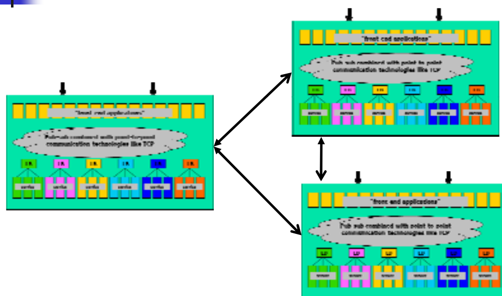

More technical barriers

- Most data centers are interconnected by
  - Extremely fast links (10-40 Gbit)
  - But with high latency
- Protocols such as TCP can't run at high speeds unless latency is *low*
- This implies that we may need new protocols if we plan to interconnect data centers over large scale

Understanding Trends

- Basically two options
  - Study the fundamentals
  - Then apply to specific tools
- Or
  - Study specific tools
  - Extract fundamental insights from examples

Understanding Trends

- Basically two options
  - Study the fundamentals
  - Then apply to specific tools
- Or
  - Study specific tools
  - Extract fundamental insights from examples

Ken's bias

- I work on reliable, secure distributed computing
  - Air traffic control systems, stock exchanges, electric power grid
  - Military "Information Grid" systems
  - Modern data centers
- To me, the question is:

  How can we build systems that do what we need them to do, reliably, accurately, and securely?

Butler Lampson's Insight

- Why computer scientists didn't invent the web
  - CS researchers would have wanted it to "work"
  - The web doesn't really work
  - But it doesn't really need to!
- Gives some reason to suspect that Ken's bias isn't widely shared!

Example: Air Traffic Control using Web technologies

- Assume a "private" network
- Web browser could easily show planes, natural for controller interactions
- What "properties" would the system need?
  - Clearly need to know that trajectory and flight data is current and consistent
  - We expect it to give sensible advice on routing options (e.g. not propose dangerous routes)
  - Continuous availability is vital: zero downtime
    - Expect a soft form of real-time responsiveness
  - Security and privacy also required (post 9/11!)

ATC systems divide country up

France

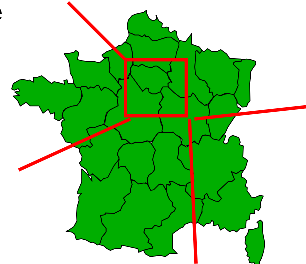

More details on ATC

- Each sector has a control center
- Centers may have few or many (50) controllers
  - In USA, controller works alone
  - In France, a "controller" is a team of 3-5 people
- Data comes from a radar system that broadcasts updates every 10 seconds
- Database keeps other flight data
- Controllers each "own" smaller sub-sectors

Issues with old systems

- Overloaded computers that often crash
  - Attempt to build a replacement system failed, expensively, back in 1994
- Getting slow as volume of air traffic rises
- Inconsistent displays a problem: phantom planes, missing planes, stale information
- Some major outages recently (and some near-miss stories associated with them)
  - TCAS saved the day: collision avoidance system of last resort… and it works….

Concept of IBM's 1994 system

- Replace video terminals with workstations
- Build a highly available real-time system guaranteeing no more than 3 seconds downtime per year
- Offer much better user interface to ATC controllers, with intelligent course recommendations and warnings about future course changes that will be needed

ATC Architecture

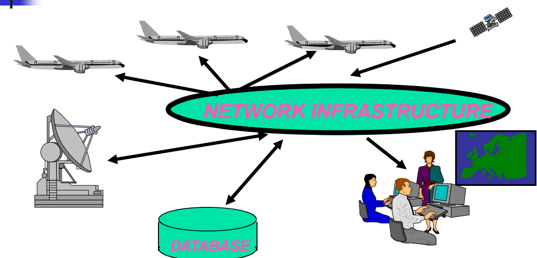

So… how to build it?

- In fact IBM project was just one of two at the time; the French had one too
  - IBM approach was based on lock-step replication
    - Replace every major component of the system with a fault-tolerant component set
    - Replicate entire programs ("state machine" approach)
  - French approach used replication selectively
    - As needed, replicate specific data items.
    - Program "hosts" a data replica but isn't itself replicated

IBM: Independent consoles… backed by ultra-reliable components

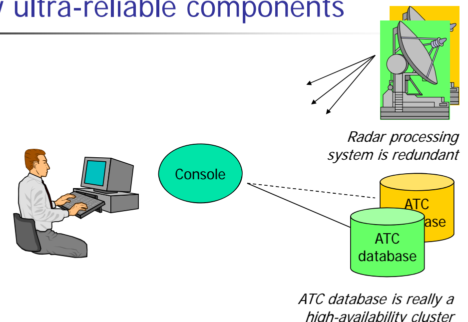

France: Multiple consoles… but in some ways they function like one

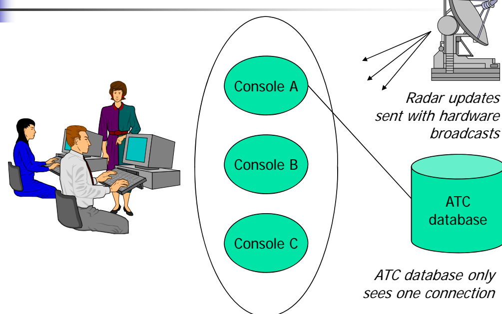

Different emphasis

- IBM imagined pipelines of processing with replication used throughout. "Services" did much of the work.

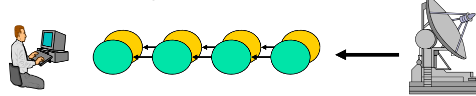

- French imagined selectively replicated data, for example "list of planes currently in sector A.17"
  - E.g. controller interface programs could maintain replicas of certain data structures or variables with system-wide value
  - Programs did computing on their own helped by databases

Other technologies used

- Both used standard off-the-shelf workstations (easier to maintain, upgrade, manage)
  - IBM proposed their own software for fault-tolerance and consistent system implementation
  - French used Isis software developed at Cornell
- Both developed fancy graphical user interface much like the Web, pop-up menus for control decisions, etc.

IBM Project Was a Fiasco!!

- IBM was unable to implement their fault-tolerant software architecture! Problem was much harder than they expected.
  - Even a non-distributed interface turned out to be very hard, major delays, scaled back goals
  - And performance of the replication scheme turned out to be terrible for reasons they didn't anticipate
- The French project was a success and never even missed a deadline… In use today.

Where did IBM go wrong?

- Their software "worked" correctly
  - The replication mechanism wasn't flawed, although it was much slower than expected
- But somehow it didn't fit into a comfortable development methodology
  - Developers need to find a good match between their goals and the tools they use
  - IBM never reached this point
- The French approach matched a more standard way of developing applications

ATC problem lingers in USA…

- "Free flight" is the next step
  - Planes use GPS receivers to track own location accurately
  - Combine radar and a shared database to see each other
  - Each pilot makes own routing decisions
  - ATC controllers only act in emergencies
- Already in limited use for long-distance flights

Free Flight (cont)

- Now each plane is like an ATC workstation
- Each pilot must make decisions consistent with those of other pilots
  - … but if FAA's project failed in 1994, why should free flight succeed in 2010?
  - Something is wrong with the distributed systems infrastructure if we can't build such things!
- In CS514, learn to look at technical choices and steer away from high-risk options

Impact of technology trends

- Web Services architecture should make it much easier to build distributed systems
  - Higher productivity because languages like Java and C# and environments like J2EE and .NET offer powerful help to developers
- The easy development route inspires many kinds of projects, some rather "sensitive"
  - But the "strong" requirements are an issue
    - Web Services aren't aimed at such concerns

Examples of mission-critical applications

- Banking, stock markets, stock brokerages
- Heath care, hospital automation
- Control of power plants, electric grid
- Telecommunications infrastructure
- Electronic commerce and electronic cash on the Web (very important emerging area)
- Corporate "information" base: a company's memory of decisions, technologies, strategy
- Military command, control, intelligence systems

We depend on distributed systems!

- If these critical systems don't work
  - When we need them
  - Correctly
  - Fast enough
  - Securely and privately
- … then revenue, health and safety, and national security may be at risk!

Critical Needs of Critical Applications

- Fault-tolerance: many flavors
  - Availability: System is continuously "up"
  - Recoverability: Can restart failed components
- Consistency:
  - Actions at different locations are consistent with each other.
  - Sometimes use term "single system image"
- Automated self-management
- Security, privacy, etc….:
  - Vital, but not our topic in this course

So what makes it hard?

- ATC example illustrated a core issue
- Existing platforms
  - Lack automated management features
  - Handle errors in ad-hoc, inconsistent ways
  - Offer one form of fault-tolerance mechanism (transactions), and it isn't compatible with high availability
- Developers often forced to step outside of the box… and might stumble.
  - But why don't platforms standardize such things?

End-to-End argument

- Commonly cited as a justification for *not* tackling reliability in "low levels" of a platform
- Originally posed in the Internet:
  - Suppose an IP packet will take n hops to its destination, and can be lost with probability p on each hop
  - Now, say that we want to transfer a file of k records that each fit in one IP (or UDP) packet
  - Should we use a retransmission protocol running "end-to-end" or n TCP protocols in a chain?

End-to-End argument

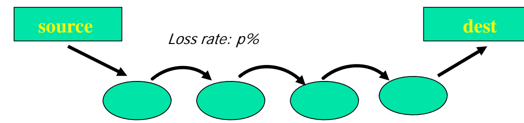

Probability of successful transit: $(1-p)^n$,
Expected packets lost: $k-k*(1-p)^n$

Saltzer et. al. analysis

- If p is very small, then even with many hops most packets will get through
  - The overhead of using TCP protocols in the links will slow things down and won't often benefit us
  - And we'll need an end-to-end recovery mechanism "no matter what" since routers can fail, too.
- Conclusion: let the end-to-end mechanism worry about reliability

Generalized End-to-End view?

- Low-level mechanisms should focus on speed, not reliability
- The application should worry about "properties" it needs
- OK to violate the E2E philosophy if E2E mechanism would be much slower

E2E is visible in J2EE and .NET

- If something fails, these technologies report timeouts
  - But they also report timeouts when nothing has failed
  - And when they report timeouts, they don't tell you what failed
  - And they don't offer much help to fix things up after the failure, either

Example: Server replication

- Suppose that our ATC needs a highly available server.
- One option: "primary/backup"
  - We run two servers on separate platforms
  - The primary sends a log to the backup
  - If primary crashes, the backup soon catches up and can take over

Split brain Syndrome…

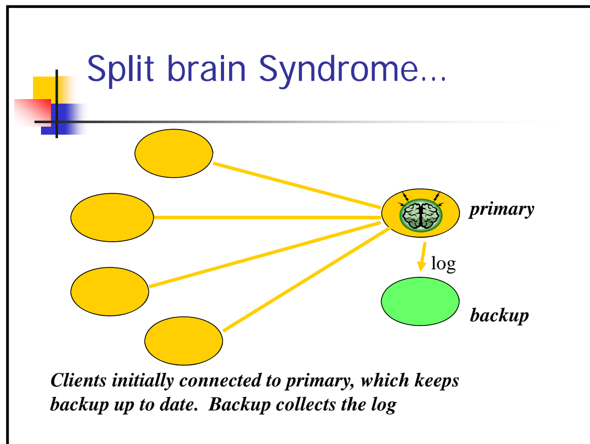

Clients initially connected to primary, which keeps backup up to date. Backup collects the log

Split brain Syndrome…

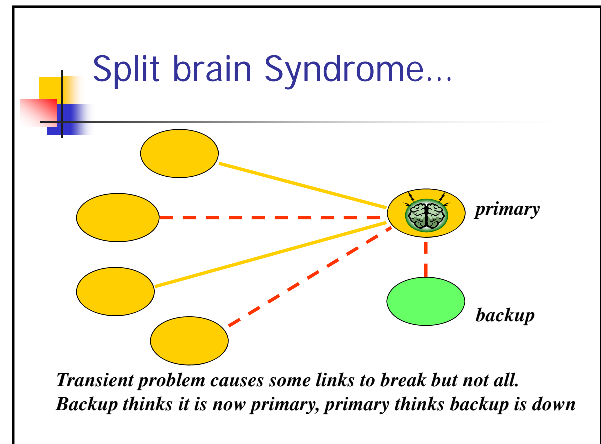

Transient problem causes some links to break but not all. Backup thinks it is now primary, primary thinks backup is down

Split brain Syndrome

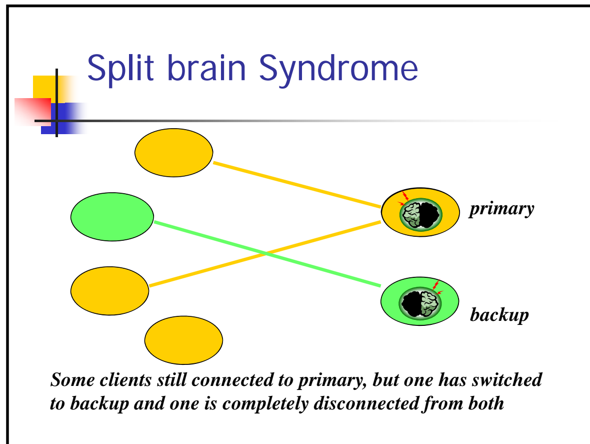

Some clients still connected to primary, but one has switched to backup and one is completely disconnected from both

Implication?

- Air Traffic System with a split brain could malfunction disastrously!
  - For example, suppose the service is used to answer the question "is anyone flying in such-and-such a sector of the sky"
  - With the split-brain version, each half might say "nope"… in response to different queries!

Can we fix this problem?

- No, if we insist on an end-to-end solution
  - We'll look at this issue later in the class
  - But the essential insight is that we need some form of "agreement" on which machines are up and which have crashed
  - Can't implement "agreement" on a purely 1-to-1 (hence, end-to-end) basis.
    - Separate decisions can always lead to inconsistency
    - So we need a "membership service"… and this is fundamentally not an end-to-end concept!

Can we fix this problem?

- Yes, many options, once we accept this
  - Just use a single server and wait for it to restart
    - This common today, but too slow for ATC
  - Give backup a way to physically "kill" the primary, e.g. unplug it
    - If backup takes over… primary shuts down
  - Or require some form of "majority vote"
    - Ad mentioned, maintains agreement on system status
- Bottom line? You need to anticipate the issue… and to implement a solution.

CS514 project

- We'll work with Web Services
  - .NET with ASP.NET in the language of your preference (C# is our favorite)
  - Or Java/J2EE
- We'll extend the platform with features like replication for high availability, self-management, etc
- And we'll use this in support of a mission critical application, mostly as a "demo"

You can work in small teams

- Either work alone at first. For third assignment can form a team of 2 or 3 members
  - Teams should tackle a more ambitious problem and will also face some tough coordination challenges
  - Experience is like working in commercial settings…

Not much homework or exams

- In fact, probably *no* graded homework or graded exams
  - But we may assign thought problems to help people master key ideas
- Grades will be based on the project
  - Can be used as an MEng project if you like
  - In this case, also sign up for CS790 credits

Textbook and readings

- Ken's textbook (came out in 2005 and already seeming a tiny bit out of date!)
  - He's planning to revise it eventually…
  - … but in distributed systems, everything is always changing!
- Additional readings: Web page has references and links