Mission-critical distributed systems

Cornell researchers in distributed systems have earned a reputation for solving difficult problems in reliable distributed applications and systems. Abstractions —and their implementations— are our product; deployment in real systems is how their impact is measured. By that measure, our impact has been substantial. For example, many of the building blocks used in today's distributed systems can trace their roots back to research in our department, including the fail-stop processor abstraction, fault-tolerant broadcast, state machine replication, virtual synchrony, and failure detectors.

Cornell-developed systems that employ these abstractions are widespread. One example is CS faculty Ken Birman's Isis Toolkit, which runs the New York and Swiss Stock Exchanges, the French Air Traffic Control System, and the US Navy's AEGIS warship. Also, the CORBA fault-tolerance standard is based on Isis-style process-group replication.

The later Horus and Ensemble systems of Birman and CS researcher Robbert van Renesse have found their way into several well-known systems. IBM's Websphere product employs a version of Ensemble to implement high availability, and Microsoft's forthcoming Windows Cluster technology draws on ideas and an architecture first demonstrated in Horus. CS professor Fred Schneider and van Renesse are working with an intranet search engine company to incorporate their "chain replication" algorithm into products running on many large and well-known Web sites.

CS professor Gün Sirer's work examines the construction of large-scale, resilient infrastructure services based on the newly emerging "peer-to-peer" paradigm, in which clients act as both resource consumers and providers, greatly improving the performance and security properties of the system. Sirer's work on the CoDoNS system has shown how to build a peer-to-peer name service for the Internet and has been deployed to serve the Internet domain name space for all of China.

Cornell's strength in this area can be traced to theoretical and practical computer scientists working side by side. Of those, Birman, Schneider, van Renesse, and Sirer remain at Cornell. Dale Skeen left to found a series of companies that specialized in reliable communication infrastructures. Ozalp Babaoglu, Keith Marzullo, Sam Toueg, and Thorsten von Eicken, who built their careers here, are credited with a wide range of seminal contributions to the theory and practice. And, after transforming himself into the world's foremost expert on scalability while at Cornell, Werner Vogels is now CTO at Amazon.com.

Distributed computing will play an increasingly critical role in the global cyber-infrastructure. The need for trustworthy systems (we use the term in the broadest sense) has received tremendous press and government attention. Massive data centers are appearing everywhere and surging in size. Programmers who used to build software for a single machine are now being asked to port their code to run on distributed platforms and to think about self-management, self-diagnosis, self-repair, and stability under stress.

We study precisely these problems at Cornell. For example, by fusing classical protocol architectures with ideas from peer-to-peer computing (such as "gossip" styles of data replication, which mimic the spread of an epidemic through a dense population), we have obtained high reliability protocols that scale seemingly without limit and offer a wide range of self-managing properties.

Ken Birman's (right) system Isis Toolkit is used in the NY and Swiss Stock Exchanges. Gün Sirer's system CoDoNS is being deployed to serve the Internet domain name space for all of China.

"My hope," says Birman, "is that by integrating these new technologies into Web services, we can offer developers turnkey high-integrity computing. A Web services developer, faced with the need to improve reliability or scalability, would push a button, fill out a property sheet, and see a clustered, scalable solution emerge."

After 25 years of developing mission-critical distributed systems, we expect to have pulled all the pieces together into a package that large numbers of developers can master and deploy. "I've worked in the field since its inception," says Birman, "and I've never felt closer to finally having the whole story and seeing the technology take off."

This story illustrates one of Cornell's enduring strengths: the ability to marshal a broad, sustained response, literally over decades, harnessing skills in both theory and practice.

TRUST

The TRUST consortium
Team for Research in Ubiquitous Secure Technology

Who trusts computers and the Internet? Every day, it seems, a credit-card database is stolen, a new worm or virus is unleashed, military computers are breached, a university computing system is broken into, a phishing scam is exposed —the list goes on and on. Even the breakdown of the Ohio power grid a few years ago was computer-related.

A new $19 million NSF multi-institutional Center has been created to research this problem of trust and security in cyberspace. TRUST researchers want to develop new technologies that will radically transform the ability of organizations to design, build, and operate trustworthy information systems that control critical infrastructure. Besides figuring out how to protect networks from attacks, they also want to develop ways for systems to "degrade gracefully" when attacked, so that the systems keep running properly. Reliability is also an important issue, because the networks we increasingly rely on are vulnerable not only to intrusion but also to breakdown. The electric power grid is a prime example and will be a test bed for the research.

The proposed research goes far beyond technical considerations; TRUST relies on collaboration with experts in economics, public policy, social science, law, and human-computer interface, for technology developed without attention to these areas risks irrelevance. Cyberspace trust is truly a global, all-encompassing problem. TRUST also has an education and outreach component, geared to K-12 schools, undergraduate students, and institutions serving underrepresented populations, which will lay the groundwork for educating new scientists and engineers to develop the next generation of trustworthy systems.

CS at Cornell is poised to take a leading role in TRUST. The Chief Scientist is CS professor Fred Schneider. Other members include CS faculty Ken Birman and Gün Sirer and ECE faculty Stephen Wicker, Rajit Manohar, and Lang Tong.

It's not surprising that Cornell is playing a critical role. Schneider chaired the National Academy study that produced *Trust in Cyberspace* and is Director of Cornell's Information Assurance Institute. Moreover, Cornell's research in mission-critical distributed systems (see the facing page) and language-based security (p. 14) will play a key role in the TRUST effort.

Academic partners	Industrial affiliates
CMU	Bellsouth
Cornell	Cisco Systems
Mills College	ESCHER
San Jose State	Hewlett Packard
Smith College	IBM
Stanford	Intel
U.C. Berkeley	Microsoft
Vanderbilt	Oak Ridge National Lab
	Qualcomm
	Sun Microsystems
	Symantec

Bob Constable and student Mike O'Donnell publish *A Programming Logic* (Winthrop).

Daniel Leivant joins.

1979

Cornell adopts the Cornell Program Synthesizer for instruction in programming. Tim Teitelbaum and student Tom Reps developed this precursor to today's integrated development environments (IDEs) for teaching a subset of PL/1 on Terak microcomputers, replacing the batch-processing punch-card system then in use. In 1980–1981, the Cornell Program Synthesizer is distributed to 80 institutions.

Gerry Salton becomes Chair of ACM SIGIR.

Bengt Aspvall, John Gilbert, Sam Toueg join.

1980

CS obtains a $2.6 million, 5-year CER (Coordinated Experimental Research) grant, a major step in increasing its presence in experimental computing.

Ozalp Babaoglu, Paul Pritchard, Dale Skeen, Tom Coleman join.

1981

David Gries publishes *The Science of Programming* (Springer-Verlag), which brings ideas on the formal development of programs to the undergrad level.

Kevin Karplus, Ken Birman join. David Gries becomes Chair. CS grows to 20 faculty.

1982

The 1982 NRC Assessment of Research-Doctorate CS programs places Cornell fifth out of 58 departments.

Bob Constable, with students Johnson and Eichenlaub, publishes a book on their verifier: *Introduction to the PL/CV Programming Logic* (Springer-Verlag).

Gerry Salton receives the first SIGIR Award for outstanding contributions to information retrieval.

Dina Bitton, Greg Johnson, Abha Moitra join.

1983

Tom Coleman publishes *Large Sparse Numerical Optimization* (Springer-Verlag LNCS 165).

CS begins to move into interdisciplinary work, helping to start a new graduate field of "manufacturing systems engineering".