BLOG @CACM

The *Communications* Web site, <http://cacm.acm.org>, features more than a dozen bloggers in the BLOG@CACM community. In each issue of *Communications*, we will publish selected posts or excerpts.

Follow us on Twitter at <http://twitter.com/blogCACM>

DOI:10.1145/2811284

<http://cacm.acm.org/blogs/blog-cacm>

The Morality of Online War; the Fates of Data Analytics, HPC

John Arquilla considers justifications for warfare in the cyber realm, while Daniel Reed looks ahead at big data and exascale computing.

John Arquilla
"The Ethics of Cyberwar"

<http://bit.ly/1LFEU2g>

July 2, 2015

All over the world, there is a growing sense conflict is spreading from the physical realm to the virtual domain. The 2007 cyberattacks on Estonia, the military use of cyberwar techniques in the 2008 Russo-Georgian War, and the "cybotage" committed against Iran's nuclear program by the Stuxnet (<http://bit.ly/1KMCIo0>) worm are salient signs of a growing trend. These likely form the tip of an iceberg, as cyber attacks and counterattacks can be observed in many other places. It is high time, as this new mode of conflict diffuses in breadth and deepens in intensity, to think through the ethics of cyberwar.

Under what conditions should one engage in cyberwar? How should such a conflict be waged? These questions speak to the classical division in ethical thought about warfare that addresses the matter of going from peace to war justly, then ponders how to fight one's

battles honorably. In terms of going to war justly, there are three commonly held principles: *Right purpose*, which refers mostly to acting in self-defense; *Due authority* seeks authorization from a national or supranational body; and *Last resort*, which is self-explanatory. Ideas of fighting justly cluster around *Noncombatant immunity*, a focus on military vs. civilian targets, and *Proportionality*, avoiding excessive force.

Right purpose has always been a fraught element of just-war theory and practice. As Napoleon once said, "I had to conquer Europe to defend France." Many military adventures follow similar logic, justifying acts of aggression as preemptive or preventive defensive actions. Stuxnet would fall in the ethically dodgy area of prevention, and one can see how cyber attack may move nations in the direction of preemptive and preventive action. Not good.

Due authority, until the Information Age, was confined to nations, coalitions, or even transnational bodies like the United Nations. NATO made choices to intervene militarily in Kosovo in

1999, and in recent years in Libya. The U.N. authorized action to repel invading North Korean forces in 1950; and so on. This category includes and allows ethical choices to go to war made by individual nations—even when that choice might have been made in error (like the U.S.-led war against Iraq in 2003, whose justification was the mistaken belief Saddam Hussein had, or soon would have, weapons of mass destruction). In cyberwar, "due authority" suffers because armies, navies, and air forces are not necessary; just malicious software and skilled hackers. "Authority" loses meaning in a world where aggressive networks, or even highly adept individuals, can wage cyberwar.

Last resort typically has referred to a requirement to pursue diplomatic efforts until it is clear they will not resolve a given crisis. This aspect of just-war theory has also proved a bit nebulous, as sometimes war is resorted to because one or another party to a dispute just gets tired of negotiating. The July Crisis of 1914 that led to World War I falls in this category. The Japanese-American talks in 1941 were frustrating enough to Tokyo that the choice was made to attack Pearl Harbor before diplomatic talks ended. When it comes to cyberwar, its fundamentally covert, deniable nature may mean it will be used *during* negotiations—clearly the case with Stuxnet.

Noncombatant immunity is the principle to avoid deliberate targeting of civilians. Over the past century, it has been outflanked by technologies that allow the innocent to be struck directly, without prior need to defeat armed forces.

es protecting them. World War II saw deliberate burning of many cities—and nuclear attacks on civilians in Japan as soon as the atomic bomb became available. During the Korean War, virtually every building in Pyongyang was flattened, and a greater weight of bombs fell on North Vietnam in “the American War” than were dropped on Hitler’s Germany. How will this principle play out in an era of cyberwar? With far less lethal harm done to noncombatants, but no doubt with great economic costs inflicted upon the innocent.

Proportionality has proved less difficult to parse over the past century or so. By and large, nuclear-armed nations have refrained from using ultimate weapons in wars against others not so armed. Korea stayed a conventional conflict; Vietnam, too, even though the outcomes of both for the nuclear-armed U.S. were, in the former case an uneasy draw, in the latter an outright defeat. In cyberwar, the principle of proportionality may play out more in the type of action taken, rather than in the degree of intensity of the action. A cyber counterattack in retaliation for a prior cyber attack generally will fall under the proportionality rubric. When might a cyber attack be answered with a physically destructive military action? The U.S. and Russia have both elucidated policies suggesting they might respond to a “sufficiently serious” cyber attack by other-than-cyber means.

Classical ideas about waging war remain relevant to strategic and policy discourses on cyberwar. Yet, it is clear conflict in and from the virtual domain should impel us to think in new ways about these principles. In terms of whether to go to war, the prospects may prove troubling, as cyber capabilities may encourage preemptive action and erode the notion of “war” as a tool of last resort. When it comes to strictures against targeting civilians (so often violated in traditional war), cyberwar may provide a means of causing disruption without killing many (perhaps not any) civilians. Yet there are other problems, as when non-state actors outflank the “authority” principle, and when nations might employ disproportionate physical force in response to virtual attack.

In 1899, when advances in weapons technologies made leaders wary of the costs and dangers of war, a conference (<http://bit.ly/1KMCJZg>) was held at The

Hague to codify the ethics and laws of armed conflict, followed by another meeting on the same subject in 1907. Perhaps it is time to go to The Hague again, as a new realm of virtual conflict has emerged. Even if we cannot live up to ethical ideals that might be agreed upon in such a gathering, it is imperative the world community should make the effort. Now.

Daniel A. Reed
“Exascale Computing
and Big Data:
Time to Reunite”

<http://bit.ly/1SQ0X8w>
 June 25, 2015

In other contexts, I have written about cultural and technical divergence of the data analytics (also known as machine learning and big data) and high-performance computing (big iron) communities. I have called them “twins separated at both” (in <http://bit.ly/1M186kd> and <http://bit.ly/1IUkOSF>). They share technical DNA and innate behaviors despite superficial differences. After all, they were once united by their use of BSD UNIX and SUN workstations for software development.

Both have built scalable infrastructures using high-performance, low-cost x86 hardware and a suite of (mostly) open source software tools. Both have addressed ecosystem deficiencies by developing special-purpose software libraries and tools (such as SLURM (<http://bit.ly/1M18i32>) and Zookeeper (<http://bit.ly/1IUl3xl>) for resource management and MPI (<http://bit.ly/1E4Ij41>) and Hadoop (<http://bit.ly/1IHHR1b>) for parallelism), and both have optimized hardware for problem domains (Open Compute (<http://bit.ly/1DlipOT>) for hardware building block standardization, FPGAs (<http://bit.ly/1KMEFRs>) for search and machine learning, and GPU accelerators for computational science).

I have seen this evolution in both the HPC and cloud computing worlds. One reason I went to Microsoft was to bring HPC ideas and applications to cloud computing. At Microsoft, I led a research team (<http://bit.ly/1K179nC>) to explore energy-efficient cloud hardware designs and programming models, and I launched a public-private partnership between Microsoft and the National Science Foundation on cloud applications (<http://bit.ly/1hfZr1V>). Back in aca-

demia, I seek to bring cloud computing ideas to HPC.

Jack Dongarra and I co-authored an article for *Communications* on the twin ecosystems of HPC and big data and the challenges facing both. The article (<http://bit.ly/1If45X0>) examines commonalities and differences, and discusses unresolved issues associated with resilience, programmability, scalability, and post-Dennard hardware futures (<http://bit.ly/1Dlj1E3>). The article makes a plea for hardware and software integration and cultural convergence.

The possibilities for this convergence are legion. The algorithms underlying deep machine learning (<http://bit.ly/1gEXlsr>) would benefit from parallelization and data movement minimization techniques commonly used in HPC applications and libraries. Similarly, approaches to failure tolerance and systemic resilience common in cloud software have broad applicability to high-performance computing. Both domains face growing energy constraints on the maximum size of systems, necessitating shared focus on domain-specific architectural optimizations that maximize operations per joule.

There is increasing overlap of application domains. New scientific instruments and sensors produce unprecedented volumes of observational data, and intelligent *in situ* algorithms are increasingly required to reduce raw data and identify important phenomena in real time. Conversely, client-plus-cloud services are increasingly model-based, with rich physics, image processing, and context that depend on parallel algorithms to meet real-time needs.

The growth of Docker (<http://bit.ly/1IHIHLI>) and containerized (<http://bit.ly/1DljqGL>) software management speaks to the need for lightweight, flexible software configuration management for increasingly complex software environments. I hope we can develop a unified hardware/software ecosystem leveraging the strengths of each community; each would benefit from the experiences and insights of the other. It is past time for the twins to have a family reunion. ■

John Arquilla is a professor at the U.S. Naval Postgraduate School. **Daniel A. Reed** is Vice President for Research and Economic Development, University Computational Science and Bioinformatics Chair, and professor of Computer Science, Electrical and Computer Engineering, and Medicine at the University of Iowa.

© 2015 ACM 0001-0782/15/10 \$15.00