The Man-Machine Integration Era

Daniel Bilar[1]

[1]**Norwich University, 158 Harmon Drive, Northfield Vermont 05663, USA**

Corresponding author:
Daniel Bilar[1]

Email address: dbilar@norwich.edu

ABSTRACT

Through sensor ubiquity, the man-machine integration era is upon us. This integration is taking place in distributed, continuously changing, optimizing/learning, finite precision feedback systems. Security challenges of such systems abound, also due to constantly co-evolving threat actors and changing environments.
Are we adequately preparing to defend such systems, and more importantly: how do we ensure they are worth defending? This position paper posits we are ill-prepared, due to perverse incentives affecting methodology, results and foundational corpuses.
With respect to the first question, we will corroborate this ill-preparedness in the context of a basic requirement – situational awareness – for so-called Moving Target Defenses. We'll argue that the second question hinges on a deceptively straightforward permanence invariant: The ability to robustly encode the infinite value of a human being in finite precision systems. Here too, we are failing to develop needed toolsets and skills.
Successfully tackling both questions may determine the future winning model of the man-machine integration era; whether we'll lose the Republic of the People to the People's Republic.

MOTIVATION

> En 1975 et en France, il y a des urgences. Et me voici, plongé dans mon époque, dans mon pays et dans leur quotidian. Si donc la meilleure part de moi-même chante une hymne sourde, une autre part vit l'aujourd'hui détestable où se débat et s'enfonce mon pays. Et mon Europe. Et mon Occident. Et s'il est bien de proposer une morale, il faut aussi proposer un combat et un terrain pour celui-ci. La morale, je l'ai dite et chantée. Le combat, il est celui d'Occident
>
> Jean Cau (1975) "Pourquoi La France" [1]

Parce qu'il y a des urgences. Some years ago, we obtained through private channels what plausibly looked like a Chinese/Taiwanese Information Warfare (IW) curriculum. Part of it contained the standard lesson plan in network security, circuit design, compilers, operating systems, cryptography, software engineering; then branched out into multi-media technology, mathematical modelling, EM leaks, computer virus design, emergency response, numerical analysis, and more. Beyond the breadth and depth, what struck non-military CS educators such as ourselves was a multi-domain systemic focus: US and Taiwanese social information systems, campaign science, synthetic experiments, hacker methods, laws and regulations, command systems, Revolution in Military Affairs, C4I and more.

This curriculum reflected an approach that is deeply holistic across scales and domains; "attentive", in Tony Corn's memorable phraseology, "to the rhythm of civilizations and the chronopolitical dimension of statecraft" [2]. We are ignoring these dynamics to our detriment, which is why this position paper was written.

THE MAN-MACHINE INTEGRATION ERA (MMIE)

What manner of machine marks our era? A machine that maps man to finite numbers through sensors instrumenting private and common spaces: personal cell-phones, home energy; public transport,

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

communication, consumption, finance, city infrastructure, eyes in the sky. These numbers are used to drive decisions such as credit scores and insurance rates, air flight watch lists, prison sentencing and probation guidelines[1], as well as partner and investment recommendations, health care allocations, and more.

This machine is overwhelmingly driven by corporate entities. The high-speed, high-availability requirements of leaders such as Google, Amazon and Netflix impel technical innovations at an astounding rate: Since 2002, Amazon Web Services (AWS) has made world-wide distributed utility computing available with per-use pricing; Netflix innovated multiple evolutions of cloud-based distributed architecture, continuous delivery pipelines, instrumentation analysis technologies for self-optimization and fault injection frameworks ("Chaos Engineering"[2]) for their world-wide content delivery systems. It took Google 22 days (!) (from tested silica to data center) to deploy custom ASICs with "reduced computational precision" managing power consumption. These innovations are made available to the public at very affordable prices, sometimes advertised as 'free' – the true cost being paid with user behaviour profiling (and subsequent mapping to finite numbers).

Distributed, concurrently executing dynamic systems pose particular headaches, both at the single system thread-interleaved and wider distributed system level. A ten-year retrospective (2005-2014) of 145 papers on debugging concurrent and multicore systems identified major gaps, among them items such as validation, evaluation and metrics [3]. Distributed systems (our focus here) subsume single-system problems and add on: Certain bug classes are unique to such systems (such as data consistency, scalability and topology bugs) and some bugs are 'killer bugs' which can cascade and immobilize multiple nodes, or the entire cluster. A careful 'cloud bug study' (11 person/year effort) was undertaken to classify and manually annotate thousands of issues in six popular distributed systems (*Hadoop MapReduce, HDFS, HBase, Cassandra, Zookeeper*, and *Flume*) along multiple dimensions. The study found that every implication, from failed operation, to performance degradation, downtime, data loss, data corruption, loss and staleness can be caused by virtually any software and hardware fault combination [4].

More recently, the same group homed in on 104 distributed concurrency (DC) bugs (bugs triggered by unexpected timing of events). They created a detailed multi-dimensional taxonomy; analyzed timing issues, trigger pre-conditions, error and fixing strategies. Their findings are striking and well worth studying, among them: We lack tools to analyze complex protocol interactions. Distributed model checkers have triggering blind spots due to intractability of event state space. Injecting delays at runtime seems to prevent 40% of DC bugs from triggering, but may introduce hanging risks. Almost half of DC bugs lead to silent failures, and possible mysterious errors much later in time. Fix strategies are challenging because correctly implemented synchronization of globally distributed systems is a hard problem [5]. We understand predictably as of yet little about the event timings and hardware/software constellations which violate the implicit and explicit system assumptions.

Moreover, incipient shifts in the operating environment – notably the long-coming transition to a global IPv6 address space (2^{128} vs 2^{64} addresses in IPv4) – pose additional difficulties. Modern scanners like Zmap can scan the entire IPv4 space in under an hour. Scanning the IPv6 address space in this fashion would take 10^{22} years. Such scaling barriers preclude many IPv4 defense schemes, including URL content categorization, IP reputation systems and IP blacklisting [6]. Adapting darknets (routable address spaces with no active services except passive packet collectors) for IPv6 space in the wider context of IPv6 situational awareness remains an active research area [7].

Dangers to the Republic

We will argue that the current trajectory of the MMIE is working against the United States in at least two ways: First, it disproportionately increases threat vectors affecting US societal stability; secondly, it is at odds with US traditions on the inviolable rights of the individual.

We explore our position by sketching the unaddressed challenges to imbue the MMIE with safeguards against optimizing the individual away. Of perhaps more pressing concern, we are also failing to work towards the needed technical defenses of MMIE systems. As we will argue, this is due to structural problems in (cyber-security) science. We start by discussing system defense lacunae by illustrating self-same in the area of so-called *Moving Target Defenses (MTDs)*

[1] https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing
[2] http://techblog.netflix.com/2014/09/introducing-chaos-engineering.html

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

MOVING TARGET DEFENSES

The National Science Council gave an early MTD definition in 2011: "The concept of controlling change across multiple system dimensions in order to increase uncertainty and apparent complexity for attackers."

One such example is 'port hopping', where assignments of (TCP/UDP) ports to network accessible services change more or less rapidly. The idea is to decrease the information value of prior service port reconnaissance efforts vis-a-vis the attackers (see for taxonomies of mechanism, patterns and recent research [8]). Operationally, MTDs have been discussed to defend Smart Grids and ICS/SCADA systems [9, 10].

One salient point to note is that MTDs change a running, working system while needing to preserve at the very least mission-sustaining operational availability. The good guys should not be unduly affected, for otherwise, engaging in such defenses constitutes a 'self-DoS', a self-inflicted denial of service. For MTD schemes to be operationally practical, they must provide stability guarantees. At the very least predictable oscillations, better stable on average and preferably strict. Exceptional situations notwithstanding (eg the 2001 Hainan Island US Navy EP-3E signal plane incident[11]) engaging in MTDs should not lead to self-defeating irrecoverable instabilities.

In addition, MTD mechanisms seek to deny to an attacker a 'true' (current, useful) view of pertinent system state variables (such as port to service mappings), whilst maintaining good guy's mission capabilities. MTDs have information theoretical aspects: Minimizing useful information leaks to the attacker (including leaks arising from MTD interactions with the environment, attacker and side channels) whilst maintaining 'truer' situational awareness. By 'truer' we mean that the defender must be able to 'learn' (ascertain) the state of pertinent variables faster than the attacker, either weakly on average or strongly at all times. The latter point dovetails with framing of MTDs as a defensive form of asymmetric cyber-space operations [12].

It should surprise no one that figuring out stability and information theoretical aspects of such embedded MTDs is non-trivial. Abstracted analytic models illuminate general directions and serve as theoretical sanity checks. Some such sanity checks are well known (Rice's theorem), some are known to smaller communities (impossibility predicting actions of rational agents [13]), some are virtually unknown (impossibility of leak elimination [14]), and some are waiting to be rediscovered (distributed systems theorems [15]). We maintain that due to the complexities of unfolding open system dynamics, the royal road to operations runs through scientific experiments on real-life systems and/or large-scale simulations. In the ideal case, performing such scientific experiments over time amasses robust results, yielding useful predictions and perhaps even unexpected cross-domain insights [16].

Presently, it would be operationally *irresponsible* to deploy MTD schemes because, among other things, we cannot meet a fundamental MTD pre-condition: the ability to accurately quantify situational awareness. We cannot meet this pre-condition because we either lack or cannot trust the empirical data, methods and results in the existing literature.

Malaise of (Cyber-security) Science

Robust science has a few characteristics, among them reproducibility, repeatability, and internal as well as external validity [17]. Due to misaligned incentive structures and poor methodology, much of scientific literature is suffused with 'results' that cannot be repeated by the same team, let alone be reproduced by others.

1. Incentives: A mean *bon mot* states that"poor methods get results"; publishable results to be precise, in high impact journals, cited by others. 'Number of publications' is the crude coin of the researcher realm, used by academic and governmental decision makers for promotions, grant decisions, continuing employment, intra-field glory, and mutual support. Reviews by peers and editorial discretion are meant to enact quality control.

   Significant problems have been reported with this setup. There is a publication bias towards positive results considered 'significant' (typically by a so-called null-hypothesis test with a cut-off of 1-5%). This bias makes negative results (i.e. Edison's 'I know 10,000 things that don't work') almost impossible to publish in venues that matter. To a lesser extent, this is also true of replication studies. These structural failures incentivize researchers at best to focus on quick-turnaround 'least-publishable unit' (LPU) research, fishing for positive results with resulting poor

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

validity (eg ‘p-hacking’), and at worst to commit fraud by peer review collusion[3] and by faking data, clumsily [18].

2. Methods: A recent Nature survey of 1,500+ researchers found that 70% tried and failed to reproduce another scientist’s experiments, and more than half have failed to reproduce their own experiments [19]. This is consistent with the worsening unfolding reproducibility crisis in much of science, including sociology, psychology, medicine, and biology (see Fig. 1).

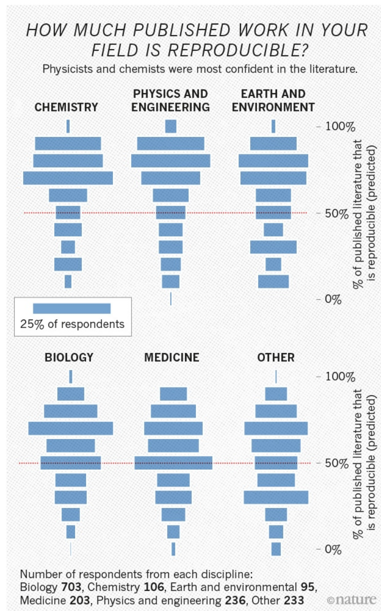

Figure 1. Research reproducibility in science. Graphic from [19]. Reproduced with permission.

In biotechnology, more than half of studies published in reputed journals (such as Science, Nature, Cell, PNAS) are reportedly not reproducible by industrial labs [20]. Tim Horton, the editor of the Lancet (a leading medical journal) put it thusly:

> The case against science is straightforward: much of the scientific literature, perhaps half, may simply be untrue. Afflicted by studies with small sample sizes, tiny effects, invalid exploratory analyses, and flagrant conflicts of interest, together with an obsession for pursuing fashionable trends of dubious importance, science has taken a turn towards darkness.

In some fields, an estimated 85% of research resources are thus wasted [21].

Security Science and MTD

Cybersecurity is not exempt from the general science malaise, which partly explains the lack of a foundational corpus. In an exemplary but all too rare meta-study [22], a researcher reviewed ninety cybersecurity papers between 1981 and 2008 and evaluated their security perspective, target of quantification, underlying assumptions and type of validation:

> The result shows how the validity of most methods is still strikingly unclear. Despite applying a number of techniques from fields such as computer science, economics and reliability theory to the problem it is unclear what valid results exist with respect to operational security.

[3] http://retractionwatch.com/2014/07/08/sage-publications-busts-peer-review-and-citation-ring-60-
?utm_source=dlvr.it&utm_medium=twitter

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

Of the 166 papers produced in 2015 under NSA Research Directorate's Science of Security (SoS) initiative , security metrics yielded about 20 papers; less than a handful treated MTDs explicitly, and only one seems to have dealt with gathering a large-scale empirical data corpus (in the context of software patch deployment). Aggregate results have yet to gel into methodical prescriptions and best practices (notable exception [23]), though SoS has announced that they plan to expand their outreach and partnerships with industry [24].

Without taking this argument too far (right at the line is Herley's epistemological unfalsifiability of security claims [25]), our main point is that there are foundational scientific challenges that need to be resolved first before any results can be translated to the operational arena.

The NSA SoS initiative has taken up the task to put (cyber) security science on a more solid foundation. They focus on "Five Hard Problems":

1. Scalability and Composability
2. Policy-Governed Secure Collaboration
3. Security Metrics Driven Evaluation, Design, Development, and Deployment
4. Resilient Architectures
5. Understanding and Accounting for Human Behavior

Relevant advances include empirical studies on threat landscape changes post-security tech introduction, formal mathematical frameworks for precise specification of security properties (among them availability, usability, scalability, evolvability, and resilience), and practical mathematical frameworks supporting reasoning about cyber-physical system robustness. Other researchers have suggested the sensible point that cybernetic (control) principles should inform cyber-security science, with explicit calls for guidance from military science [26].

PERMANENCE IN FINITE SYSTEMS

MMIE systems are built to continuously improve. Since no human can manage thousands or even hundreds of interacting parameters, optimization and tuning of system parameters is increasingly left to machine learning and applied AI [27, 28] . Though the numeric range of digital finite precision systems seems daunting, notions of infinity are handled poorly (in IEEE754 by defining a very large and very small value as $\pm\infty$). A very large number (say 1.1897×10^{4932} in quadruple precision 128 bit IEEE754 format) is still not qualitatively different from 27 in the Cantor Hierarchy of Infinity sense.

Since a defining characteristic of the MMIE is the mapping of man to finite numbers, a particular set of unaddressed challenges we face is thus: How do we ensure the permanence, the indelibility, the infinite value of human beings as invariants such that the AI decision procedures will not optimize us away?

Amodei et al explores unintended harmful behaviour that may arise in real world AI systems, if the designers did not anticipate certain failure modes. One such failure mode to avoid is 'reward hacking', eg gaming the programmed goal linked to a reward function. In the case of a cleaning robot with the goal of a mess-free office, reward hacking may take the form of disabling its vision (so as not see messes), or eliminating causes of messes (throwing out all movable objects). Several solution directions are offered for simple cases. Complicated reward functions over longer time scales are still in uncharted territory [29]. Horvitz, in a recent June talk at CMU, issued best practices for safe AI that dovetail with Amadei at al's concerns, among them disclosure of parameters on failure rates, tradeoffs and preferences, and transparency of perception, inference, and action [30].

Transparency as to which entities are optimized across which scales and domains is going to be crucial for the human value invariant in light of reward hacking. Thompson's fascinating 1996 experiment serves as an early cautionary tale [31]. His goal was to use genetic algorithms (GA, a set of optimization methods) to evolve a 10*10 cell circuit on a 64*64 cell FPGA (a configurable chip with cells consisting of transistors) that could distinguish between a 1kHz and a 10 kHz sound wave. The circuit was unclocked, meaning that the GA was not evolving a digital system, but an analog continuous-time dynamical system of transistors (with input period five orders of magnitude longer than input to output signal propagation time). The solution the GA found after 2-3 weeks had surprisingly properties:

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

Certain FPGA cells outside the 10*10 solution circuit with no connected wire path to influence the circuit could not be removed without negatively affecting the solution. This meant that the GA included unexpected properties of the FPGA physical substrate, EM coupling or the power supply in its search space.

The solution was as a result also non-transferrable to other 10*10 cell patches, nor other nominally identical FGPAs with normal manufacturing variation. One may be inclined to dismiss physical properties leveraged by software optimization as far removed from today's systems. This would be ill-advised, as the Rowhammer DRAM memory cell flip attacks (from JavaScript!) demonstrate [32]. It does not strain credulity to imagine AI 'reward hacking' MMIE systems (in conjunction with opaque signals) leading to different outcomes in a testing or simulation environment than in operationally settings.

Assuming we can solve the issue of infinity as applied permanence invariant in an optimization-resistant manner, the decision problems run deeper. A recast of Asimov's Laws would be insufficiently expressive as a basis for AI decision ethics. Such a view classifies ethical actions as forbidden, obligatory and morally indifferent. Selmer Bringsjord (a cognitive science and computer science professor at RPI) shows that ethics based solely on Deontic laws are "painfully naïve and morally inexpressive", lacking for example superogatory (good to do, but not forbidden) and suberogatory (bad to do, but not forbidden) considerations. He and his lab have been working on richer computational logical L_{EH} embedded in his 21st century Ethical Hierarchy (EH) for almost a decade [33].

We'll treat this topic more extensively in future work. In our view, negotiating this admittedly unusual terrain in the MMIE should be a matter of national security.

HUMANE STABILITY

What do shortcomings in cyber security science, in understanding MMIE systems, and the mapping of man to finite numbers have to do with our national security and China? We maintain that the MMIE, in its current trajectory, works again the United States in two ways, by

1. Increasingly building accessible threat vectors to US societal stability

   Unpalatable consequences of the MMIE attack surface were starkly laid bare by the Center for Long-Term Cybersecurity at UC Berkeley. They published an "imaginative map of possibility space", describing five eminently plausible cyber security scenarios we'll find ourselves in come 2020. Arguably, the first one "The New Normal" already materialized, with the 2nd "Omega" and 5th scenarios "Internet of Emotions" likely consequences of sensor ubiquity and man-machine integration era [34]. It is likely that this state of affairs endangers the United States more than China because of the deeper and wider networked computerization of our infrastructure, both critical and civilian.

2. Increasingly clashing with US values regarding inalienable rights of the individual

   Events in China's modern history demonstrate that the focus of the Chinese Communist Party (CCP) lay in the collective, with the individual more of a *quantité négligeable*. Examples include Mao's Great Leap Forward and the man-made, calculated famine that starved to death upwards of 40 million human beings in the early 1960s [35]. For the ambitious Three Gorges Dam energy project, the CCP initiated the forced relocation of more than a million people, with hundreds of communities destroyed. In contradistinction, since its inception, the US has championed negative rights, and put a premium on the inviolable sphere of the individual. We are in danger of abdicating our moral high ground by failing to embed the concept of the inviolable individual into the MMIE.

We posit that the CCP has a long-term vision of the man-machine integration era: monitored, collectivized, and controlled, with the hard constraint of societal stability. Julien's masterly exposition of the Chinese 勢 (Shi) helps explore this premise [36]. In 勢, reality is perceived as an arrangement of things to be relied upon and worked to one's advantage.勢's characteristics are formal, dynamic, strategic; one of its metaphors is the womb. Our closest conception would be setting up pieces for the long game (see also [37]); in aseptic DARPA-speak, 勢 would correspond to creating "the drivers towards convergence".

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

One such driver is the CCP's 网格化管理 ("grid management") system. Though CCP surveillance is not new [38], this modern data analytics based city population surveillance system aims to monitor for reported signs of instability and quell incipient social unrest [39]. A complimentary nation-wide "Social Credit System" will be mandatory by 2020. This system establishes a citizen's score between 350-950, with sensor data drawn from shopping, credit, and social networks giving insight into public activities, private consumption, individual habits, and the behaviour of family and friends [40]. These scores may determine private and government perks such as financing rates, travel permits, perhaps elite university admissions. It is not a stretch to conclude that social mobility may be linked to political compliance. Already in 2013, the Supreme People's Court established a blacklist database. This blacklist (shared with China Securities Regulatory Commission and the Credit Reference Center at the People's Bank of China) serves to keep convicted debtors from luxury purchases, such as high-end cars and train tickets [41].

Marzak et al's analysis of China's "Great Firewall" (GFW) and "Great Cannon" (GC) exposes another set of such convergence drivers [42]. We'd classify the GFW as a 'man-on-the-loop' versus the GC's 'man-in-the-loop' method. The GFW acts as a stateful, top-down content policy enforcer, terminating undesirable TCP connections *in situ* by injecting forged TCP reset segments. Individual segments are reassembled prior to being passed on to the decision logic for improved censorship granularity. The GC trades whole stream reassembly for down-selected target addresses. It hijacks unencrypted TCP connections (so-called man-in-the-middle attack) and injects malicious content against perceived antagonists (such as GreatFire, a GFW watchdog). As Fig. 2 illustrates, constitutive elements of anti-censorship networks such as Tor are targeted, as well.

Figure 2. How the GFW works. Drawing by Christian Zenker [43], reproduced under CC01 1.0 license with permission.

The CCP is arranging to birth a future. We are not paying enough attention.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

THINKING IN SYSTEMS

Suggestions for Education

We are failing to focus on the skills needed to defend and imbue our values into the distributed systems of the MMIE. We maintain this not on the basis of a thorough content and methods assessment of the 20 top CS programs (with disproportionate influence on CS education as a whole [44]), but from our personal experiences and interactions with graduates of such programs, academic/DoD researchers, Program Managers and DoD solicitations. In this assessment, we are bolstered by David Brumley, a noted cybersecurity professor at Carnegie Mellon University with one of the very best US programs:

> We agree with the overarching problem: the U.S. is not doing enough to train students in cybersecurity and universities must do more to grow their cybersecurity curricula. We are not alone: A number of institutions — like Polytechnic Institute of NYU, U.C. Santa Barbara, Berkeley, and many others — are also helping to address the problem through action: building courses that create cybersecurity experts.

In Table 1, we give a brief non-exhaustive list of skills and exemplar reference projects as starting points. In ABET parlance, we envision a learning outcomes syncretism evinced by performance indicators such as Figs. 3 and 4. A detailed curriculum toward this end is planned as future work.

Skills	Projects
Experimental Design	CMU 15-321: Research Methods for Experimental Computer Science [45]
Information Theory	MICMINE: Maximal information-based nonparametric exploration statistics [46]
Control and Optimization	Remy: Computer-generated, adaptive TCP congestion control algorithms [47]
Modeling and Invariants	Eureqa/Nutonian: Evolutionary search for invariant discovery [48]
Exact Software Engineering and Verification	eARF: Architectural reasoning framework from design to requirements [49]
Simulation	Robots that adapt like animals [50]
Representation and Precision	UNUM: Precise, error-bounded number format suited for parallelism [51]
Envisioning Information	Visualization of high-dimensional complex data [52]
Mechanism Design	DARPA Red Balloon [53]
Reasoning with Infinities	Omega++: Coq-certified decision procedures for Presburger arithmetic with infinity [54]
Applied Theoretic Computer Science	Parsimony: Construction of small Turing Machines independent of ZFC [55]

Table 1. Skills for the MMIE.

Not listed but of considerable importance are research incentive reforms to meliorate the state of published scientific research. A remediation blueprint from an early warning voice, John Ioannidis, lists appropriate incentives and drivers: Changes to academic and grant reward systems, the adoption of a replication culture, better study designs, reproducibility practices, large scale collaboration among other things "to make more published research true" [56]. We close our discussion with suggestions on ways forwards anent MTDs.

Suggestions for Operational MTD

MTDs function in operational environments that are in flux, and as such are subject to adversarial dynamics. Instrumentation and measurement must cover several domains, among them the MTD itself, defended system, the operating environment and adversaries. Space and scope constraints do not afford a deeper treatment; we note pars pro toto that anent IPv6, such research has barely begun [57]. In addition, we need to capture measurement metadata, which at the very minimum would

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

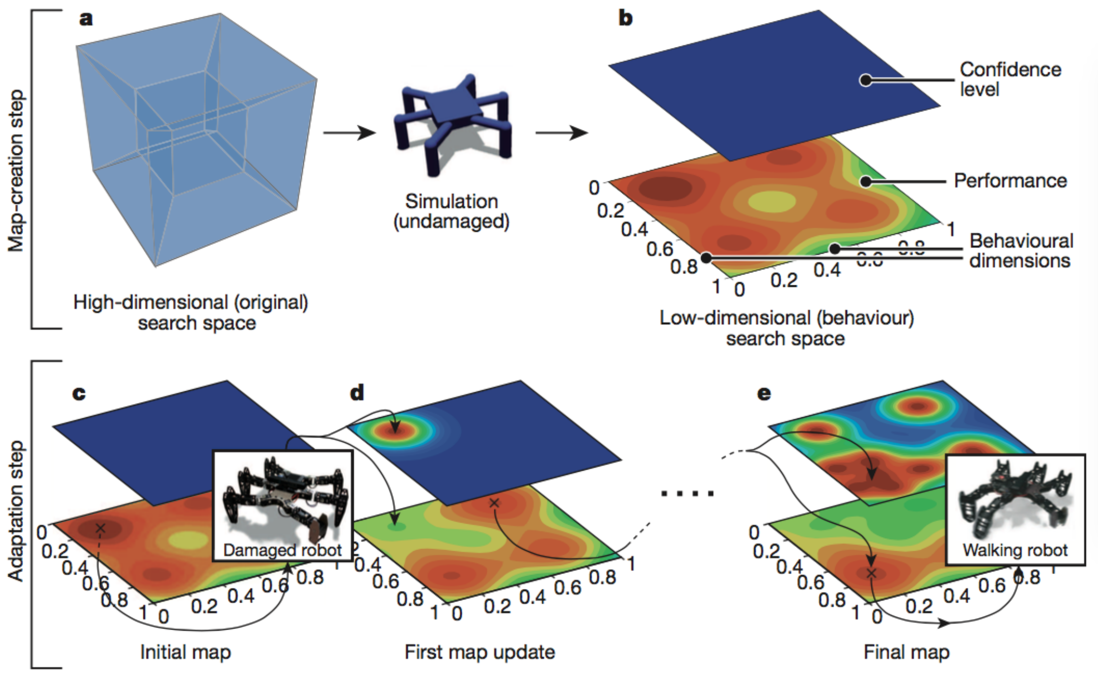

Figure 3. Evidence of outcome syncretism: Creating a "behavior-performance" map, illustrating simulation/optimization/visualization skills. High-dimensional space is searched to find performance potential, including high-performing behavior at each point in low-dim behavior space. Figure from [50] reproduced with permission.

include a detailed source (provenance) description, a quality / validity model (life cycle/decay rate), as well as traditional uncertainty indicators such as standard error and confidence intervals when taking the measurement. For prediction, we need to perform quantifiable experiments under solid multi-factorial regimes with wide coverage [58]. Ideally, we'd explain phenomenological results by means of an appropriate generative model to hedge against spurious correlation fishing.

Since operational MTD by its nature introduces changes with system-wide information dissemination requirements (eg if you mutate service mapping to confuse the bad guys, the good guys still need to know the new mappings), the MTD modeling should be informed by concepts addressing these issues. Specifically, we recommend framing:

1. MTD as distributed (control) systems, with focus on the trade off between information consistencies versus availability properties. This is also known as the "uniformity of information" problem (FLP/CAP [59]). Doyle's mathematical control architecture framework with its explicit treatment of design space trade-off of multiple metrics tackles this head-on [60].There also have been theoretical stabs at investigating latency-aware algorithms for decisions at runtime [61].

2. MTD as targets, with focus on neuralgic attack surfaces. It is an open question to what extent proposed MTD dynamic network schemes [62] can be resilient to denial and degradation attacks when their control systems are targeted [63].

3. MTD as asymmetric adversarial learning, with focus on information leaks and emitted side channels. We offer [64] as a starting point to quantitatively assess general information leakage bounds robust with respect to operational scenarios.

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

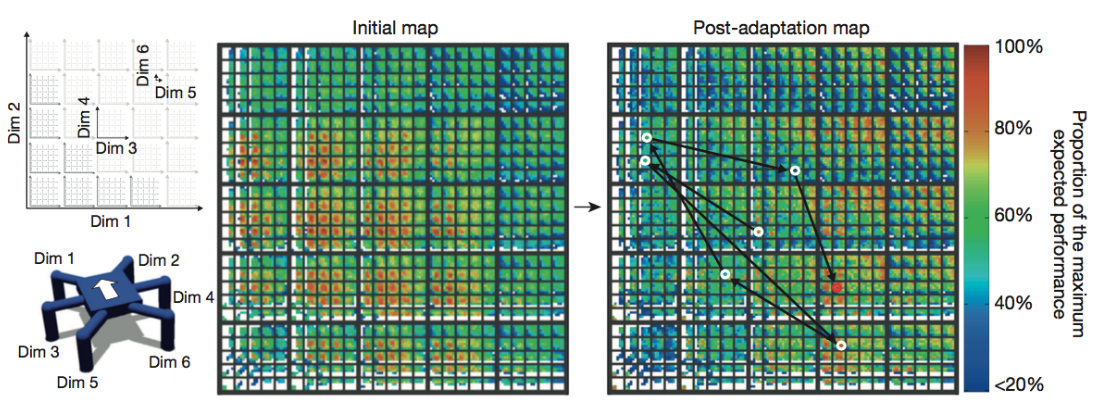

Figure 4. 'Behavior-performance" map storing high-performing behaviors at each point in a six-dimensional behavior space. Each colored pixel represents the highest-performing behavior discovered during map creation. Matrices visualize the six-dimensional behavioral space in two dimensions according to the legend in the top-left. Created with the Open Dynamics Engine physics simulator (http://www.ode.org). Figure from [50] reproduced with permission.

EPILOGUE

> Blessings are not found in something that has been weighed, nor in something that has been measured, nor in something that has been counted, only in something that is hidden from the eyes.
>
> Talmud, Bava Metziah 42a

Whereas the authoritarian CCP is implementing societal instrumentation, feedback and control systems to dampen instabilities, we have yet to demonstrate any pressing awareness, let alone articulate a compelling American counter-vision for the man-machine integration era. On the contrary, the "Disuniting of America", of which US historian Arthur Schlesinger warned 25 years ago [65] is progressing, with all the instabilities such societal fragmentation and polarization entails. One may smirk at some of the more puerile manifestations, such as calls from CEOs and 'thinkers' in the Bay Area tech segment to secede from non-tech hoi polloi. But more serious among the casualties is the eroding notion of sacrosanct individual rights: Judge Ruth Bader Ginsburgh expressing abroad her preferences for positive rights charters such Canada's or South Africa's over the US' as international model. 27% of US college students assent to restricting offensive political views on campus . A recent Gallup poll July 3rd, 2016 records a new low: Only 52% are "extremely proud" to be American. These are worrying symptoms, consistent with *Zivilisationsmüdigkeit*.

It is thus not inconceivable that US citizens may select a model for the man-machine era at odds with our history; one that values utilitarian, finitely quantified measures over deontological individual rights endowed by a Creator. We may lose the Republic of the People to the People's Republic. In different time, in a different America, a similar warning was already given, and better, by Chinese dissident Wei Jingsheng. His testimony concludes the transcript of the 106th Congress regarding China's accession to the WTO [66]:

> The U.S. should recognize that after the fall of the Soviet Union, Communist China is democracy's most formidable adversary [..]If the United States will not fight the world's largest tyranny politically, then inevitably, it will have to fight it economically, and eventually, militarily. Therefore, the only way to preserve peace and freedom begins by comprehending democracy's greatest enemy, and countering it effectively.

ACKNOWLEDGMENTS

We would like to thank Adam Elkus (GMU) for bringing "Robot Ethics" and the Defense Science Board report on Autonomy June 2016 to our attention. We acknowledge Hamed Okhravi (MIT LL) point about "predictable oscillations between points" with respect to MTD stability. 'Thinking in Systems' is a homage to the unsurpassed eponymous systems classic by Donella Meadows. We thank Elisabeth Bilar for repeatedly reviewing this manuscript and clarifying language, crystallizing argument and organization.

REFERENCES

[1] J. Cau, *Pourquoi la France*. Paris: La Table Ronde, 1975.

[2] T. Corn, "World War IV as fourth-generation warfare," Tech. Rep., 2006. [Online]. Available: http://www.academia.edu/download/35827462/World{_}War{_}IV{_}as{_}Fourth-Generation{_}Warfare.pdf

[3] S. Asadollah, D. Sundmark, S. Eldh, and H. Hansson, "10 Years of research on debugging concurrent and multicore software: a systematic mapping study," *Software Quality*, 2016. [Online]. Available: http://link.springer.com/article/10.1007/s11219-015-9301-7

[4] H. S. Gunawi, V. Martin, A. D. Satria, M. Hao, T. Leesatapornwongsa, T. Patana-anake, T. Do, J. Adityatama, K. J. Eliazar, A. Laksono, and J. F. Lukman, "What Bugs Live in the Cloud?" in *Proceedings of the ACM Symposium on Cloud Computing - SOCC '14*. New York, New York, USA: ACM Press, 2014, pp. 1–14. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2670979.2670986

[5] T. Leesatapornwongsa, J. Lukman, and S. Lu, "TaxDC: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems," *and Operating Systems*, 2016. [Online]. Available: http://dl.acm.org/citation.cfm?id=2872374

[6] Q. Li, C. Larsen, and T. van der Horst, "IPv6: A Catalyst and Evasion Tool for Botnets and Malware Delivery Networks," *Computer*, vol. 46, no. 5, pp. 76–82, 5 2013. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6297968

[7] L. Hendriks, A. Sperotto, and A. Pras, "Characterizing the IPv6 security landscape by large-scale measurements," in *IFIP International Conference on*, 2015. [Online]. Available: http://link.springer.com/chapter/10.1007/978-3-319-20034-7{_}16

[8] G. Cai, B. Wang, Y. Luo, S. Li, and X. Wang, "Characterizing the running patterns of moving target defense mechanisms," in *2016 18th International Conference on Advanced Communication Technology (ICACT)*. IEEE, 1 2016, pp. 191–196. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7423324

[9] S. Groat, M. Dunlop, and W. Urbanksi, "Using an IPv6 moving target defense to protect the Smart Grid," *2012 IEEE PES*, 2012. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs{_}all.jsp?arnumber=6175633

[10] C. Davidson and T. Andel, "Feasibility of Applying Moving Target Defensive Techniques in a SCADA System," in *International Conference on Cyber*, 2016. [Online]. Available: http://search.proquest.com/openview/ffef4f3ba393afed6debfa50a581b0cc/1?pq-origsite=gscholar{&}cbl=396500

[11] S. Kan, R. Best, C. Bolkcom, and R. Chapman, "China-US Aircraft Collision Incident of April 2001: Assessments and Policy Implications," *CRS Report for*, 2001. [Online]. Available: http://www.intelligencelaw.com/files/pdf/law{_}library/crs/RL30946{_}10-10-2001.pdf

[12] A. Sumari and D. Gunawan, "Cyberspace Operations as Multiplier Power in Asymmetric Conflict," in *Conference on Cyber Warfare and Security*, 2014. [Online]. Available: https://www.academia.edu/23465126/Cyberspace{_}Operations{_}as{_}Multiplier{_}Power{_}in{_}Asymmetric{_}Conflict

[13] D. Foster and H. Young, "On the impossibility of predicting the behavior of rational agents," *of the National Academy of Sciences*, 2001. [Online]. Available: http://www.pnas.org/content/98/22/12848.short

[14] J. J. Harmsen and W. A. Pearlman, "Capacity of Steganographic Channels," *IEEE Transactions on Information Theory*, vol. 55, no. 4, pp. 1775–1792, 4 2009. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4802306

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

[15] F. Fich and E. Ruppert, “Hundreds of impossibility results for distributed computing,” *Distributed Computing*, vol. 16, no. 2-3, pp. 121–163, 9 2003. [Online]. Available: http://link.springer.com/10.1007/s00446-003-0091-y

[16] T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf, “Undecidability of the spectral gap,” *Nature*, vol. 528, no. 7581, pp. 207–211, 12 2015. [Online]. Available: http://www.nature.com/doifinder/10.1038/nature16059

[17] R. Maxion, “Making Experiments Dependable.” Springer Berlin Heidelberg, 2011, pp. 344–357. [Online]. Available: http://link.springer.com/10.1007/978-3-642-24541-1{_}26

[18] N. Brown and J. Heathers, “The GRIM test,” *PeerJ*, vol. 53, no. 9, pp. 1689–1699, 2016. [Online]. Available: https://peerj.com/preprints/2064/

[19] M. Baker, “1,500 scientists lift the lid on reproducibility,” *Nature*, vol. 533, no. 7604, pp. 452–454, 2016. [Online]. Available: www.nature.com/doifinder/10.1038/533452a

[20] B. Booth, “Academic bias and biotech failures : Trade Secrets,” 2011. [Online]. Available: http://blogs.nature.com/tradesecrets/2011/04/22/academic-bias-and-biotech-failures

[21] I. Chalmers, P. Glasziou, A. Liberati, R. Scherer, P. Langenberg, E. v. Elm, D. Cooksey, N. I. f. H. Research, S. Garattini, I. Chalmers, C. Gross, G. Anderson, N. Powe, D. Stuckler, L. King, H. Robinson, M. McKee, P. Perel, J. Miranda, Z. Ortiz, J. Casas, S. Oliver, J. Gray, D. Tallon, J. Chard, P. Dieppe, I. Chalmers, C. Rounding, K. Lock, S. Hewlett, M. D. Wit, P. Richards, e. al., N. Cooper, D. Jones, A. Sutton, N. Patsopoulos, A. Analatos, J. Ioannidis, S. Mallett, M. Clarke, C. Hewitt, S. Hahn, D. Torgerson, J. Watson, J. Bland, A. Rutjes, J. Reitsma, M. D. Nisio, N. Smidt, J. v. Rijn, P. Bossuyt, I. Chalmers, S. Hopewell, M. Clarke, L. Stewart, J. Tierney, S. Hopewell, K. Dickersin, M. Clarke, A. Oxman, K. Loudon, K. Dwan, D. Altman, J. Arnaiz, e. al., P. Glasziou, E. Meats, C. Heneghan, S. Shepperd, C. Young, and R. Horton, “Avoidable waste in the production and reporting of research evidence.” *Lancet (London, England)*, vol. 374, no. 9683, pp. 86–9, 7 2009. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/19525005

[22] V. Verendel, “Quantified security is a weak hypothesis,” in *Proceedings of the 2009 workshop on New security paradigms workshop - NSPW ’09*. New York, New York, USA: ACM Press, 2009, p. 37. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1719030.1719036

[23] J. Dykstra, *Essential Cybersecurity Science: Build, Test, and Evaluate Secure Systems*. Oreilly & Associates Inc, 2016. [Online]. Available: https://books.google.com/books?hl=en{&}lr={&}id=ExsoCwAAQBAJ{&}oi=fnd{&}pg=PT72{&}dq=related:9RD3pyiHAcwJ:scholar.google.com/{&}ots=ToHgrBUn5k{&}sig=u-W2d2X3uuuP1-TJSKBQAWWvSEo

[24] NSA, “Science of Security,” Tech. Rep., 2015. [Online]. Available: http://cps-vo.org/node/25098

[25] C. Herley, “Unfalsifiability of security claims,” *Proceedings of the National Academy of*, 2016. [Online]. Available: http://www.pnas.org/content/113/23/6415.short

[26] A. Roque, K. B. Bush, and C. Degni, “Security is about control,” in *Proceedings of the Symposium and Bootcamp on the Science of Security - HotSos ’16*. New York, New York, USA: ACM Press, 2016, pp. 17–24. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2898375.2898379

[27] M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, and N. de Freitas, “Learning to learn by gradient descent by gradient descent,” 6 2016. [Online]. Available: http://arxiv.org/abs/1606.04474

[28] R. S. Olson and J. H. Moore, “TPOT: A Tree-based Pipeline Optimization Tool for Automating Machine Learning,” in *JMLR*, 2016, p. 8. [Online]. Available: https://docs.google.com/viewer?a=v{&}pid=sites{&}srcid=ZGVmYXVsdGRvbWFpbnxhdXRvbWwyMDE2fGd4OmFmYjMyNWU2NWI1YTBmZg

[29] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, “Concrete Problems in AI Safety,” 6 2016. [Online]. Available: http://arxiv.org/abs/1606.06565

[30] E. Horvitz, “Reflections on Safety and Artificial Intelligence,” Pittsburgh, PA, 2016.

[31] A. Thompson, “An evolved circuit, intrinsic in silicon, entwined with physics.” Springer Berlin Heidelberg, 1997, pp. 390–405. [Online]. Available: http://link.springer.com/10.1007/3-540-63173-9{_}61

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

[32] D. Gruss, C. Maurice, and S. Mangard, "Rowhammer.js: A Remote Software-Induced Fault Attack in JavaScript," *arXiv:1507.06955v1*, vol. 2016, 2015. [Online]. Available: http://arxiv.org/abs/1507.06955

[33] S. Bringsjord, "A 21st-Century Ethical Hierarchy for Robots and Persons: EH." [Online]. Available: http://kryten.mm.rpi.edu/SELPAP/MOREMORALROBOTS/SBringsjord{_}ethical{_}hierarchy{_}062215b.pdf

[34] Center for Long-Term Cybersecurity, "Cybersecurity futures 2020," uc berkeley, Tech. Rep., 2016. [Online]. Available: https://cltc.berkeley.edu/scenarios/

[35] J. Yang, Edward Friedman, Jian Guo, and Stacy Mosher, *Tombstone: The great Chinese famine, 1958-1962*. New York: Farrar, Straus and Giroux, 2012.

[36] F. Jullien, *The Propensity of Things : Towards a History of Efficacy in China*. MIT Press, 1999.

[37] M. Pillsbury, *The hundred-year marathon: China's secret strategy to replace America as the global superpower*. Henry Holt and Co, 2015. [Online]. Available: https://books.google.com/books?hl=en{&}lr={&}id=grGMAwAAQBAJ{&}oi=fnd{&}pg=PT7{&}ots=Ljg6oR{_}m5u{&}sig=QzocIF9pSQp9TTiOaxYCoL{_}k0jE

[38] R. Creemers, M. Meissner, P. Crossley, P. Mattis, and S. Hoffman, "Is Big Data Increasing Beijing's Capacity for Control?" *ChinaFile Conservation*, 8 2016. [Online]. Available: https://www.chinafile.com/conversation/Is-Big-Data-Increasing-Beijing-Capacity-Control{%}3F

[39] Stratfor, "China Intensifies Its Domestic Surveillance Program — Stratfor," Tech. Rep., 2016. [Online]. Available: https://www.stratfor.com/image/china-intensifies-its-domestic-surveillance-program?id=be1ddd5371{&}uuid=bb5be779-4f41-421b-8f2a-a56e5689449f

[40] R. Creemers, "Planning Outline for the Construction of a Social Credit System (2014-2020) « China Copyright and Media," Beijing, p. GF No. (2014)21, 2015. [Online]. Available: https://chinacopyrightandmedia.wordpress.com/2014/06/14/planning-outline-for-the-construction-of-a-social-credit-system-2014-2020/

[41] C. Yin, "Debtors pay price for ignoring court rulings," 9 2015. [Online]. Available: http://usa.chinadaily.com.cn/epaper/2015-09/10/content{_}21844292.htm

[42] B. Marczak, N. Weaver, J. Dalek, R. Ensafi, D. Fifield, S. McKune, A. Rey, J. Scott-Railton, R. Deibert, and V. Paxson, "China's great cannon," Tech. Rep., 2015. [Online]. Available: http://www.academia.edu/download/37269796/Chinas{_}Great{_}Cannon.pdf

[43] C. Zenker, ""How the Great Firewall discovers hidden Circumvention Server"," 2015. [Online]. Available: https://twitter.com/xopn/status/681166666848419840

[44] A. Clauset, S. Arbesman, and D. B. Larremore, "Systematic inequality and hierarchy in faculty hiring networks," *Science Advances*, vol. 1, no. 1, pp. 1–6, 2015. [Online]. Available: http://advances.sciencemag.org/content/1/1/e1400005.abstract

[45] R. Maxion, "Research Methods for Experimental Computer Science." [Online]. Available: http://coolmon.ft.cs.cmu.edu/methods/about.shtml

[46] D. N. Reshef, Y. A. Reshef, H. K. Finucane, S. R. Grossman, G. McVean, P. J. Turnbaugh, E. S. Lander, M. Mitzenmacher, and P. C. Sabeti, "Detecting Novel Associations in Large Data Sets," *Science*, vol. 334, no. 6062, pp. 1518–1524, 2011. [Online]. Available: http://www.sciencemag.org/cgi/doi/10.1126/science.1205438

[47] K. Winstein and H. Balakrishnan, "TCP ex Machina: Computer-Generated Congestion Control," *Proc. ACM Conference on Communications Architectures, Protocols and Applications (SIGCOMM'13)*, pp. 123–134, 2013. [Online]. Available: http://nms.csail.mit.edu/papers/sigcomm13.pdf

[48] M. Schmidt and H. Lipson, "Distilling free-form natural laws from experimental data." *Science (New York, N.Y.)*, vol. 324, no. 5923, pp. 81–5, 4 2009. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/19342586

[49] N. Abbas, J. Andersson, and M. Iftikhar, "Rigorous Architectural Reasoning for Self-Adaptive Software Systems," *Proceedings of the 1st*, 2016. [Online]. Available: https://people.cs.kuleuven.be/{~{}}danny.weyns/papers/2016QRSA.pdf

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016

[50] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 7553, pp. 503–507, 5 2015. [Online]. Available: http://www.nature.com/doifinder/10.1038/nature14422

[51] J. L. Gustafson, *The End of Error: Unum Computing*. Taylor and Francis Ltd, 2015, no. November. [Online]. Available: https://www.crcpress.com/The-End-of-Error-Unum-Computing/Gustafson/p/book/9781482239867

[52] E. R. Tufte and Graphics Press, *Envisioning information*, 1st ed. Graphics Press, 1990. [Online]. Available: https://www.edwardtufte.com/tufte/books{_}ei

[53] J. C. Tang, M. Cebrian, N. a. Giacobe, H.-W. Kim, T. Kim, and D. B. Wickert, "Reflecting on the DARPA Red Balloon Challenge," *Communications of the ACM*, vol. 54, no. 4, p. 78, 2011.

[54] A. Sharma, S. Wang, A. Costea, A. Hobor, and W.-N. Chin, "Certified Reasoning with Infinity," in *FM 2015: Formal Methods: 20th International Symposium, Oslo, Norway, June 24-26, 2015, Proceedings*, N. Bjørner and F. de Boer, Eds. Cham: Springer International Publishing, 2015, pp. 496–513. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-19249-9{_}31

[55] A. Yedidia and S. Aaronson, "A Relatively Small Turing Machine Whose Behavior Is Independent of Set Theory," *arXiv preprint arXiv:1605.04343*, 2016. [Online]. Available: http://arxiv.org/abs/1605.04343

[56] J. P. A. Ioannidis, K. Boyack, R. Klavans, and et al., "How to Make More Published Research True," *PLoS Medicine*, vol. 11, no. 10, p. e1001747, 10 2014. [Online]. Available: http://dx.plos.org/10.1371/journal.pmed.1001747

[57] Caida, "Exploring the evolution of IPv6: topology, performance, and traffic," 2016. [Online]. Available: https://www.caida.org/funding/nets-ipv6/nets-ipv6{_}proposal.xml

[58] K. Killourhy and R. Maxion, "Should security researchers experiment more and draw more inferences?" *usenix.org*. [Online]. Available: https://www.usenix.org/legacy/events/cset11/tech/final{_}files/Killourhy.pdf

[59] S. Gilbert and N. Lynch, "Perspectives on the CAP Theorem," *Computer*, vol. 45, no. 2, pp. 30–36, 2 2012. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6122006

[60] N. Matni, Y. P. Leong, Y. S. Wang, S. You, M. B. Horowitz, and J. C. Doyle, "Resilience in Large Scale Distributed Systems," *Procedia Computer Science*, vol. 28, pp. 285–293, 2014. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S1877050914000994

[61] J. Cámara, G. Moreno, and D. Garlan, "Analyzing latency-aware self-adaptation using stochastic games and simulations," *Transactions on Autonomous and Adaptive Systems*, 2016. [Online]. Available: http://dl.acm.org/citation.cfm?id=2774222https://pdfs.semanticscholar.org/e466/63d8b0de44f1f181f77ec98289d3c06fdee1.pdf

[62] H. Okhravi, M. Rabe, T. Mayberry, and W. Leonard, "Survey of cyber moving target techniques," Tech. Rep., 2013. [Online]. Available: http://oai.dtic.mil/oai/oai?verb=getRecord{&}metadataPrefix=html{&}identifier=ADA591804

[63] D. Bilar, "Degradation and Subversion through Subsystem Attacks," *IEEE Security & Privacy Magazine*, vol. 8, no. 4, pp. 70–73, 7 2010. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5523869

[64] M. Alvim, K. Chatzikokolakis, and A. Mciver, "Additive and Multiplicative Notions of Leakage, and Their Capacities," in *Proceedings of the 27th Computer Security Foundations Symposium*, 2014. [Online]. Available: http://dl.acm.org/citation.cfm?id=2708705

[65] A. M. Schlesinger, *The Disuniting of America*. New York: Norton, 1992.

[66] W. Jingsheng, "ACCESSION OF CHINA TO THE WTO," 2000. [Online]. Available: https://www.gpo.gov/fdsys/pkg/CHRG-106hhrg67832/html/CHRG-106hhrg67832.htm

PeerJ Preprints | https://doi.org/10.7287/peerj.preprints.2402v1 | CC BY 4.0 Open Access | rec: 31 Aug 2016, publ: 31 Aug 2016