

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: October 25, 2019

S. Farrell  
Trinity College Dublin  
April 23, 2019

**We're gonna need a bigger threat model**  
**draft-farrell-etm-00**

Abstract

We argue that an expanded threat model is needed for Internet protocol development as protocol endpoints can no longer be considered to be generally trustworthy for any general definition of "trustworthy."

This draft will be a submission to the DEDR IAB workshop.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 25, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4](#).e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Examples of deliberate adversarial behaviour in applications	<a href="#">4</a>
<a href="#">2.1.</a>	Malware in curated application stores . . . . .	<a href="#">4</a>
<a href="#">2.2.</a>	Virtual private networks (VPNs) . . . . .	<a href="#">4</a>
<a href="#">2.3.</a>	Compromised (home) networks . . . . .	<a href="#">4</a>
<a href="#">2.4.</a>	Web browsers . . . . .	<a href="#">5</a>
<a href="#">2.5.</a>	Web site policy deception . . . . .	<a href="#">5</a>
<a href="#">2.6.</a>	Tracking bugs in mail . . . . .	<a href="#">5</a>
<a href="#">2.7.</a>	Smart televisions . . . . .	<a href="#">6</a>
<a href="#">2.8.</a>	So-called Internet of things . . . . .	<a href="#">6</a>
<a href="#">3.</a>	Inadvertent adversarial behaviours . . . . .	<a href="#">6</a>
<a href="#">4.</a>	Possible directions for an expanded threat model . . . . .	<a href="#">7</a>
<a href="#">4.1.</a>	Develop a BCP for privacy considerations . . . . .	<a href="#">7</a>
<a href="#">4.2.</a>	Consider the user perspective . . . . .	<a href="#">7</a>
<a href="#">4.3.</a>	Consider ABuse-cases as well as use-cases . . . . .	<a href="#">7</a>
<a href="#">4.4.</a>	Re-consider protocol design "lore" . . . . .	<a href="#">8</a>
<a href="#">4.5.</a>	Isolation . . . . .	<a href="#">8</a>
<a href="#">4.6.</a>	Transparency . . . . .	<a href="#">8</a>
<a href="#">4.7.</a>	Minimise . . . . .	<a href="#">8</a>
<a href="#">4.8.</a>	Same-Origin Policy . . . . .	<a href="#">9</a>
<a href="#">4.9.</a>	Greasing . . . . .	<a href="#">9</a>
<a href="#">5.</a>	Conclusions . . . . .	<a href="#">9</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">9</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">9</a>
<a href="#">8.</a>	Acknowledgements . . . . .	<a href="#">10</a>
<a href="#">9.</a>	References . . . . .	<a href="#">10</a>
<a href="#">9.1.</a>	Informative References . . . . .	<a href="#">10</a>
<a href="#">9.2.</a>	URIs . . . . .	<a href="#">12</a>
	Author's Address . . . . .	<a href="#">13</a>

## [1.](#) Introduction

[[There's a github repo for this -- issues and PRs are welcome there.  
[<https://github.com/sftcd/etm>](https://github.com/sftcd/etm) ]]

[RFC3552], [Section 3](#) defines an "Internet Threat Model" which has been commonly used when developing Internet protocols. That assumes that "the end-systems engaging in a protocol exchange have not themselves been compromised." [RFC 3552](#) is a formal part of the IETF's process as it is also [BCP72](#).

Since [RFC 3552](#) was written, we have seen a greater emphasis on considering privacy and [[RFC6973](#)] provides privacy guidance for

Farrell

Expires October 25, 2019

[Page 2]

protocol developers. [RFC 6973](#) is not a formal BCP, but appears to have been useful for protocol developers as it is referenced by 38 later RFCs at the time of writing [1].

[BCP188](#), [[RFC7258](#)] subsequently recognised pervasive monitoring as a particular kind of attack and has also been relatively widely referenced (39 RFCs at the time of writing [2]). To date, perhaps most documents referencing [BCP188](#) have considered state-level or in-network adversaries.

In this document, we argue that we need to expand our threat model to acknowledge that today many applications are themselves rightly considered potential adversaries for at least some relevant actors. However, those (good) actors cannot in general refuse to communicate and will with non-negligible probability encounter applications that are adversarial.

We also argue that not recognising this reality causes Internet protocol designs to sometimes fail to protect the systems and users who depend on those.

Discussion related to expanding our concept of threat-model ought not (but perhaps inevitably will) involve discussion of weakening how confidentiality is provided in Internet protocols. Whilst it may superficially seem to be the case that encouraging in-network interception could help with detection of adversarial application behaviours, such a position is clearly mistaken once one notes that adding middleboxes that can themselves be adversarial cannot be a solution to the problem of possibly encountering adversarial code on the network. It is also the case that the IETF has rough consensus to provide better, and not weaker, security and privacy, and has maintained that consensus over three decades, despite repeated (and repetitive;-) debates on the topic. That consensus is represented in [[RFC2404](#)], [BCP 200](#) [[RFC1984](#)] and more latterly, the above-mentioned [BCP 188](#) as well as in the numerous RFCs referencing those works. The probability that discussion of expanding our threat model leads to a change in that rough consensus seems highly remote.

It is not clear if the IETF will reach rough consensus on a description of such an expanded threat model, but we further argue that ignoring this aspect of deployed reality cannot may not bode well for Internet protocol development.

Absent such an expanded threat model, we expect to see more of a mismatch between expectations and the deployment reality for some Internet protocols.



This internet-draft is a submission to the IAB's DEDR workshop [3] and is not intended to become an RFC.

We are saddened by, and apologise for, the somewhat dystopian impression that this document may impart - hopefully, there's a bit of hope at the end;-)

## **2. Examples of deliberate adversarial behaviour in applications**

In this section we describe a few documented examples of deliberate adversarial behaviour by applications that could affect Internet protocol development. The adversarial behaviours described below involve various kinds of attack, varying from simple fraud, to credential theft, surveillance and contributing to DDoS attacks. This is not intended to be a comprehensive nor complete survey, but to motivate us to consider deliberate adversarial behaviour by applications.

Finally, we note that while we have these examples of deliberate adversarial behaviour, there are also many examples of application developers doing their best to protect the security and privacy of their users or customers. That's just the same as the case today where we need to consider in-network actors as potential adversaries despite the many examples of network operators who do act primarily in the best interests of their users.

### **2.1. Malware in curated application stores**

Despite the best efforts of curators, so-called App-Stores frequently distribute malware of many kinds and one recent study [curated] claims that simple obfuscation enables malware to avoid detection by even sophisticated operators. Given the scale of these deployments, even a small percentage of malware-infected applications can affect a huge numbers of people.

### **2.2. Virtual private networks (VPNs)**

Virtual private networks (VPNs) are supposed to hide user traffic to various degrees depending on the particular technology chosen by the VPN provider. However, not all VPNs do what they say, some for example misrepresenting the countries in which they provide vantage points. [vpns]

### **2.3. Compromised (home) networks**

What we normally might consider network devices such as home routers do also run applications that can end up being adversarial, for example DNS and DHCP attacks from home routers, or other devices in



the home. One study [home] reports on a 2011 attack that affected 4.5 million DSL modems in Brazil. The absence of software update [RFC8240] has been a major cause of these issues and rises to the level that considering allowing this as intentional behaviour is warranted.

#### **2.4. Web browsers**

Tracking of users in order to support advertising based business models is ubiquitous on the Internet today. HTTP header fields (such as cookies) are commonly used for such tracking, as are structures within the content of HTTP responses such as links to 1x1 pixel images and (ab)use of Javascript APIs offered by browsers. [tracking]

While some people may be sanguine about this kind of tracking, others consider this behaviour unwelcome, when or if they are informed that it happens, [attitude] though the evidence here seems somewhat harder to interpret and many studies (that we have found to date) involve small numbers of users. Historically, browsers have not made this kind of tracking visible and have enabled it by default, though some recent browser versions are starting to enable visibility and blocking of some kinds of tracking. Browsers are also increasingly imposing more stringent requirements on plug-ins for varied security reasons.

#### **2.5. Web site policy deception**

Many web sites today provide some form of privacy policy and terms of service, that are known to be mostly unread. [unread] This implies that, legal fiction aside, users of those sites have not in reality agreed to the specific terms published and so users are therefore highly exposed to being exploited by web sites, for example [cambridge] is a recent well-publicised case where a service provider abused the data of 87 million users via a partnership. While many web site operators claim that they care deeply about privacy, it seems prudent to assume that some (or most?) do not in fact care about user privacy, or at least not in ways with which many of their users would agree. And of course, today's web sites are actually mostly fairly complex web applications and are no longer static sets of HTML files, so calling these "web sites" is perhaps a misnomer, but considered as web applications, it seems clear that many exist that are adversarial.

#### **2.6. Tracking bugs in mail**

Some mail user agents (MUAs) render HTML content by default (with a subset not allowing that to be turned off, perhaps particularly on mobile devices) and thus enable the same kind of adversarial tracking





seen on the web. Attempts at such intentional tracking are also seen many times per day by email users - in one study [[mailbug](#)] the authors estimated that 62% of leakage to third parties was intentional, for example if leaked data included a hash of the recipient email address.

### **2.7. Smart televisions**

There have been examples of so-called "smart" televisions spying on their owners without permission [[4](#)] and one survey of user attitudes [[smarttv](#)] found "broad agreement was that it is unacceptable for the data to be repurposed or shared" although the level of user understanding may be questionable. What is clear though is that such devices generally have not provided controls for their owners that would allow them to meaningfully make a decision as to whether or not they want to share such data.

### **2.8. So-called Internet of things**

Many so-called Internet of Things (IoT) devices ("so-called" as all devices were already things:-) have been found extremely deficient when their security and privacy aspects were analysed, for example children's toys. [[toys](#)] While in some cases this may be due to incompetence rather than being deliberately adversarial behaviour, the levels of incompetence frequently seen imply that it is valid to consider such cases as not being accidental.

## **3. Inadvertent adversarial behaviours**

Not all adversarial behaviour by applications is deliberate, some is likely due to various levels of carelessness (some quite understandable, others not) and/or due to erroneous assumptions about the environments in which those applications (now) run. We very briefly list some such cases:

- o Application abuse for command and control, for example, use of IRC or apache logs for malware command and control [[5](#)]
- o Carelessly leaky buckets [[6](#)], for example, lots of Amazon S3 leaks showing that careless admins can too easily cause application server data to become available to adversaries
- o Virtualisation exposing secrets, for example Meltdown and Spectre [[7](#)] and similar side-channels
- o Compromised badly-maintained web sites, for example that have lead to massive online databases of passwords [[8](#)]



- o Supply-chain attacks, for example the Target attack [\[9\]](#)
- o Breaches of major service providers, that many of us might have assumed would be sufficiently capable to be the best large-scale "Identity providers", for example:
  - \* 3 billion accounts: yahoo [\[10\]](#)
  - \* many millions at risk: telcos selling location data [\[11\]](#)
  - \* 50 million accounts: facebook [\[12\]](#)
  - \* 14 million accounts: verizon [\[13\]](#)
  - \* "hundreds of thousands" of accounts: google [\[14\]](#)
  - \* unknown numbers, some email content exposed: microsoft [\[15\]](#)
- o Breaches of smaller service providers: Too many to enumerate, sadly

#### **[4.](#) Possible directions for an expanded threat model**

As we believe useful conclusions in this space require community consensus, we won't offer definitive descriptions of an expanded threat model but we will call out some potential directions that could be explored at the DEDR workshop and thereafter, if there is interest in this topic.

##### **[4.1.](#) Develop a BCP for privacy considerations**

It may be time for the IETF to develop a BCP for privacy considerations, possibly starting from [\[RFC6973\]](#).

##### **[4.2.](#) Consider the user perspective**

[I-D.nottingham-for-the-users] argues that, in relevant cases where there are conflicting requirements, the "IETF considers end users as its highest priority concern." Doing so seems consistent with the expanded threat model being argued for here, so may indicate that a BCP in that space could also be useful.

##### **[4.3.](#) Consider ABuse-cases as well as use-cases**

Protocol developers and those implementing and deploying Internet technologies are typically most interested in a few specific use-cases for which they need solutions. Expanding our threat model to include adversarial application behaviours [abusecases] seems likely



to call for significant attention to be paid to potential abuses of whatever new or re-purposed technology is being considered.

#### **[4.4.](#) Re-consider protocol design "lore"**

It could be that this discussion demonstrates that it is timely to reconsider some protocol design "lore" as for example is done in [[I-D.iab-protocol-maintenance](#)]. More specifically, protocol extensibility mechanisms may inadvertently create vectors for abuse-cases, given that designers cannot fully analyse their impact at the time a new protocol is defined or standardised. One might conclude that a lack of extensibility could be a virtue for some new protocols, in contrast to earlier assumptions.

#### **[4.5.](#) Isolation**

Sophisticated users can deal with some adversarial behaviours in applications by using different instances of those applications, for example, differently configured web browsers for use in different contexts. Applications (including web browsers) and operating systems are also building in isolation via use of different processes or sandboxing. Protocol artefacts that relate to uses of such isolation mechanisms might be worth considering. To an extent, the IETF has in practice already recognised some of these issues as being in-scope, e.g. when considering the linkability issues with mechanisms such as TLS session tickets, or QUIC connection identifiers.

#### **[4.6.](#) Transparency**

Certificate transparency (CT) [[RFC6962](#)] has been an effective countermeasure for X.509 certificate mis-issuance, which used be a known application layer misbehaviour in the public web PKI. While the context in which CT operates is very constrained (essentially to the public CAs trusted by web browsers), similar approaches could be useful for other protocols or technologies.

In addition, legislative requirements such as those imposed by the GDPR for subject access to data [[16](#)] could lead to a desire to handle internal data structures and databases in ways that are reminiscent of CT, though clearly with significant authorisation being required and without the append-only nature of a CT log.

#### **[4.7.](#) Minimise**

As recommended in [[RFC6973](#)] data minimisation and additional encryption are likely to be helpful - if applications don't ever see data, or a cleartext form of data, then they should have a harder



time misbehaving. Similarly, not adding new long-term identifiers, and not exposing existing ones, would seem helpful.

#### **4.8. Same-Origin Policy**

The Same-Origin Policy (SOP) [[RFC6454](#)] perhaps already provides an example of how going beyond the [RFC 3552](#) threat model can be useful. Arguably, the existence of the SOP demonstrates that at least web browsers already consider the 3552 model as being too limited. (Clearly, differentiating between same and not-same origins implicitly assumes that some origins are not as trustworthy as others.)

#### **4.9. Greasing**

The TLS protocol [[RFC8446](#)] now supports the use of GREASE [[I-D.ietf-tls-grease](#)] as a way to mitigate on-path ossification. While this technique is not likely to prevent any deliberate misbehaviours, it may provide a proof-of-concept that network protocol mechanisms can have impact in this space, if we spend the time to try analyse the incentives of the various parties.

### **5. Conclusions**

At this stage we don't think it appropriate to claim that any strong conclusion can be reached based on the above. We do however, claim that there is a topic that could be worth discussion at the DEDR workshop and elsewhere.

### **6. Security Considerations**

This draft is all about security, and privacy.

Encryption is one of the most effective tools in countering network based attackers and will also have a role in protecting against adversarial applications. However, today many existing tools for countering adversarial applications assume they can inspect network traffic to or from potentially adversarial applications. These facts of course cause tensions (e.g. see [[RFC8404](#)]). Expanding our threat model could possibly help reduce some of those tensions, if it leads to the development of protocols that make exploitation harder or more transparent for adversarial applications.

### **7. IANA Considerations**

There are no IANA considerations.





## 8. Acknowledgements

We'll happily ack anyone who's interested enough to read and comment on this. With no implication that they agree with some or all of the above, thanks to Christian Huitema and Daniel Kahn Gillmor for comments on an earlier version of the text.

## 9. References

### 9.1. Informative References

[cambridge]

Isaak, J. and M. Hanna, "User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection", Computer 51.8 (2018): 56-59, 2018, <<https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8436400>>.

[curated]

Hammad, M., Garcia, J., and S. Malek, "A large-scale empirical study on the effects of code obfuscations on Android apps and anti-malware products", ACM International Conference on Software Engineering 2018, 2018, <[https://www.ics.uci.edu/~seal/publications/2018ICSE\\_Hammad.pdf](https://www.ics.uci.edu/~seal/publications/2018ICSE_Hammad.pdf)>.

[I-D.iab-protocol-maintenance]

Thomson, M., "The Harmful Consequences of the Robustness Principle", [draft-iab-protocol-maintenance-02](#) (work in progress), March 2019.

[I-D.ietf-tls-grease]

Benjamin, D., "Applying GREASE to TLS Extensibility", [draft-ietf-tls-grease-02](#) (work in progress), January 2019.

[I-D.nottingham-for-the-users]

Nottingham, M., "The Internet is for End Users", [draft-nottingham-for-the-users-07](#) (work in progress), March 2019.

[mailbug]

Englehardt, S., Han, J., and A. Narayanan, "I never signed up for this! Privacy implications of email tracking", Proceedings on Privacy Enhancing Technologies 2018.1 (2018): 109-126., 2018, <<https://www.degruyter.com/downloadpdf/j/popets.2018.2018.issue-1/popets-2018-0006/popets-2018-0006.pdf>>.



- [RFC1984] IAB and IESG, "IAB and IESG Statement on Cryptographic Technology and the Internet", [BCP 200](#), [RFC 1984](#), DOI 10.17487/RFC1984, August 1996, <<https://www.rfc-editor.org/info/rfc1984>>.
- [RFC2404] Madson, C. and R. Glenn, "The Use of HMAC-SHA-1-96 within ESP and AH", [RFC 2404](#), DOI 10.17487/RFC2404, November 1998, <<https://www.rfc-editor.org/info/rfc2404>>.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.
- [RFC6454] Barth, A., "The Web Origin Concept", [RFC 6454](#), DOI 10.17487/RFC6454, December 2011, <<https://www.rfc-editor.org/info/rfc6454>>.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", [RFC 6962](#), DOI 10.17487/RFC6962, June 2013, <<https://www.rfc-editor.org/info/rfc6962>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", [RFC 6973](#), DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.
- [RFC8240] Tschofenig, H. and S. Farrell, "Report from the Internet of Things Software Update (IoTSU) Workshop 2016", [RFC 8240](#), DOI 10.17487/RFC8240, September 2017, <<https://www.rfc-editor.org/info/rfc8240>>.
- [RFC8404] Moriarty, K., Ed. and A. Morton, Ed., "Effects of Pervasive Encryption on Operators", [RFC 8404](#), DOI 10.17487/RFC8404, July 2018, <<https://www.rfc-editor.org/info/rfc8404>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", [RFC 8446](#), DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.



- [smarttv] Malkin, N., Bernd, J., Johnson, M., and S. Egelman, "'What Can't Data Be Used For?' Privacy Expectations about Smart TVs in the U.S.", European Workshop on Usable Security (Euro USEC) 2018, 2018, <[https://www.ndss-symposium.org/wp-content/uploads/2018/06/eurousec2018\\_16\\_Malkin\\_paper.pdf](https://www.ndss-symposium.org/wp-content/uploads/2018/06/eurousec2018_16_Malkin_paper.pdf)>.
- [toys] Chu, G., Apthorpe, N., and N. Feamster, "Security and Privacy Analyses of Internet of Things Childrens' Toys", IEEE Internet of Things Journal 6.1 (2019): 978-985., 2019, <<https://arxiv.org/pdf/1805.02751.pdf>>.
- [tracking] Ermakova, T., Fabian, B., Bender, B., and K. Klimek, "Web Tracking-A Literature Review on the State of Research", Proceedings of the 51st Hawaii International Conference on System Sciences, 2018, <<https://scholarspace.manoa.hawaii.edu/bitstream/10125/50485/paper0598.pdf>>.
- [unread] Obar, J. and A. Oeldorf-Hirsch, "The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services", Information, Communication and Society (2018): 1-20, 2018, <<https://doi.org/10.1080/1369118X.2018.1486870>>.
- [vpns] Khan, M., DeBlasio, J., Voelker, G., Snoeren, A., Kanich, C., and N. Vallina-Rodrigue, "An empirical analysis of the commercial VPN ecosystem", ACM Internet Measurement Conference 2018 (pp. 443-456), 2018, <<https://eprints.networks.imdea.org/1886/1/imc18-final198.pdf>>.

## 9.2. URIs

- [1] <https://datatracker.ietf.org/doc/rfc6973/referencedby/>
- [2] <https://datatracker.ietf.org/doc/rfc7258/referencedby/>
- [3] <https://www.iab.org/activities/workshops/dedr-workshop/>
- [4] <https://www.welivesecurity.com/2013/11/22/lg-admits-that-its-smart-tvs-have-been-watching-users-and-transmitting-data-without-consent/>
- [5] <https://security.stackexchange.com/questions/100577/creating-botnet-cc-server-what-architecture-should-i-use-irc-http>



- [6] <https://businessinsights.bitdefender.com/worst-amazon-breaches>
- [7] <https://www.us-cert.gov/ncas/alerts/TA18-004A>
- [8] <https://haveibeenpwned.com/Passwords>
- [9] <https://www.zdnet.com/article/how-hackers-stole-millions-of-credit-card-records-from-target/>
- [10] <https://www.wired.com/story/yahoo-breach-three-billion-accounts/>
- [11] <https://www.zdnet.com/article/us-telcos-caught-selling-your-location-data-again-senator-demands-new-laws/>
- [12] <https://www.cnet.com/news/facebook-breach-affected-50-million-people/>
- [13] <https://www.zdnet.com/article/millions-verizon-customer-records-israeli-data/>
- [14] <https://www.wsj.com/articles/google-exposed-user-data-feared-repercussions-of-disclosing-to-public-1539017194>
- [15] [https://motherboard.vice.com/en\\_us/article/ywyz3x/hackers-could-read-your-hotmail-msn-outlook-microsoft-customer-support](https://motherboard.vice.com/en_us/article/ywyz3x/hackers-could-read-your-hotmail-msn-outlook-microsoft-customer-support)
- [16] <https://gdpr-info.eu/art-15-gdpr/>

Author's Address

Stephen Farrell  
Trinity College Dublin

Email: [stephen.farrell@cs.tcd.ie](mailto:stephen.farrell@cs.tcd.ie)



