COMPGA11:
Research in Information Security

Steven Murdoch
University College London

based on a course by Tony Morton

Term 2 – 2016/17

Course summary

- “To develop an understanding of what research in information security is about, how to identify a contribution, what the quality standards in scientific publications are, and to study selected technical sub-topics in depth”
- “Students will be exposed to research on information security, by reading quality technical research papers in information security”
- Understand how to **interpret, summarise and write** research (important skills for your future)
- Read some important work in the field

Aims and outcomes

- “To develop an understanding of what research in information security is about,…
  - Understand different **research approaches** and the idea of scientific method
  - Recognise if a paper follows the principles of **scientific method**
  - If not, is there a justifiable reason
    - Not all topics naturally follow the scientific method e.g. papers describing frameworks
- Be able to read and **critically review** research literature in information security

Aims and outcomes

- ...how to identify a contribution,...
  - Be able to recognise, contextualise and evaluate a **contribution** to a field of work
- ...what the quality standards in scientific publications are,...
  - Able to identify a good (or bad) piece of scientific research and explain why
  - Understand what makes a good (or bad) academic paper

Aims and outcomes

- …and to study selected technical sub-topics in depth.”
    - Be able to carry out – independently - a literature review of a chosen topic in information security

Structure of course

- Week 20 Friday (this lecture)
    - Introduction
    - Dissertation project presentations (1)
- Week 21 Friday
    - The scientific process
    - Dissertation project presentations (2)
- Weeks 22–29 Fridays, excluding weeks 25 and 30
    - Student presentations and discussion
- Week 25 Friday
    - Reading week – no lecture
- Week 30 Friday (provisional)
    - Ethics (Sasse and Courtois)

Assessment

- Two information security paper reviews (20%) – 10% each
- Presentation in class (20%)
  - You are expected to attend all presentations and be able to discuss papers
- Literature review – usually, but not required to be, on the topic for your MSc dissertation (60%)
- More details later…

Types of publication venue

- Journal
  - No presentations, no meetings, just article
- Symposium/conference
  - Published proceedings, presentation at event
- Pre-print
  - Little or no peer review, just article
- Book
  - Reviewed by publisher that it will sell, but not necessarily peer review
- Workshop
  - Presentation at event, perhaps no publication

Ranking of research

- There is a desire for an objective way to decide whether research is important
- Very difficult to do reliably but you will encounter such metrics in practice
- Mostly based around bibliometrics
    - Some legitimate reason for this
    - Though mostly because it can be processed automatically

Ranking publications

- Number of citations (per year)
- Why might this not reliably represent the importance of a paper?
- Why do people cite papers?
- How might people increase their citation count?

Ranking publication venue

- Thomson Reuters impact factor = A/B where
  - A: number of citations to articles published in previous two years
  - B: number of articles published
- Many problems with bibliometrics
- Venues do have a reputation, which is somewhat consistent

Funding for publication venue

- Reader pays (most common, e.g. IEEE S&P, CCS)
  - Pay-per-article
  - Institutional subscription
- Author pays (e.g. PLoS One)
  - Normally author’s institution pays
  - Article then made available open-access
  - Exemptions often available
- Society pays (e.g. USENIX, PoPETs)
  - Society sponsors an open access publication

Ranking researchers

- “A scientist has index *h* if *h* of his/her N_p papers have at least h citations each, and the other (N_p – h) papers have no more than *h* citations each.”
  [An index to quantify an individual's scientific research output, J. E. Hirsch]

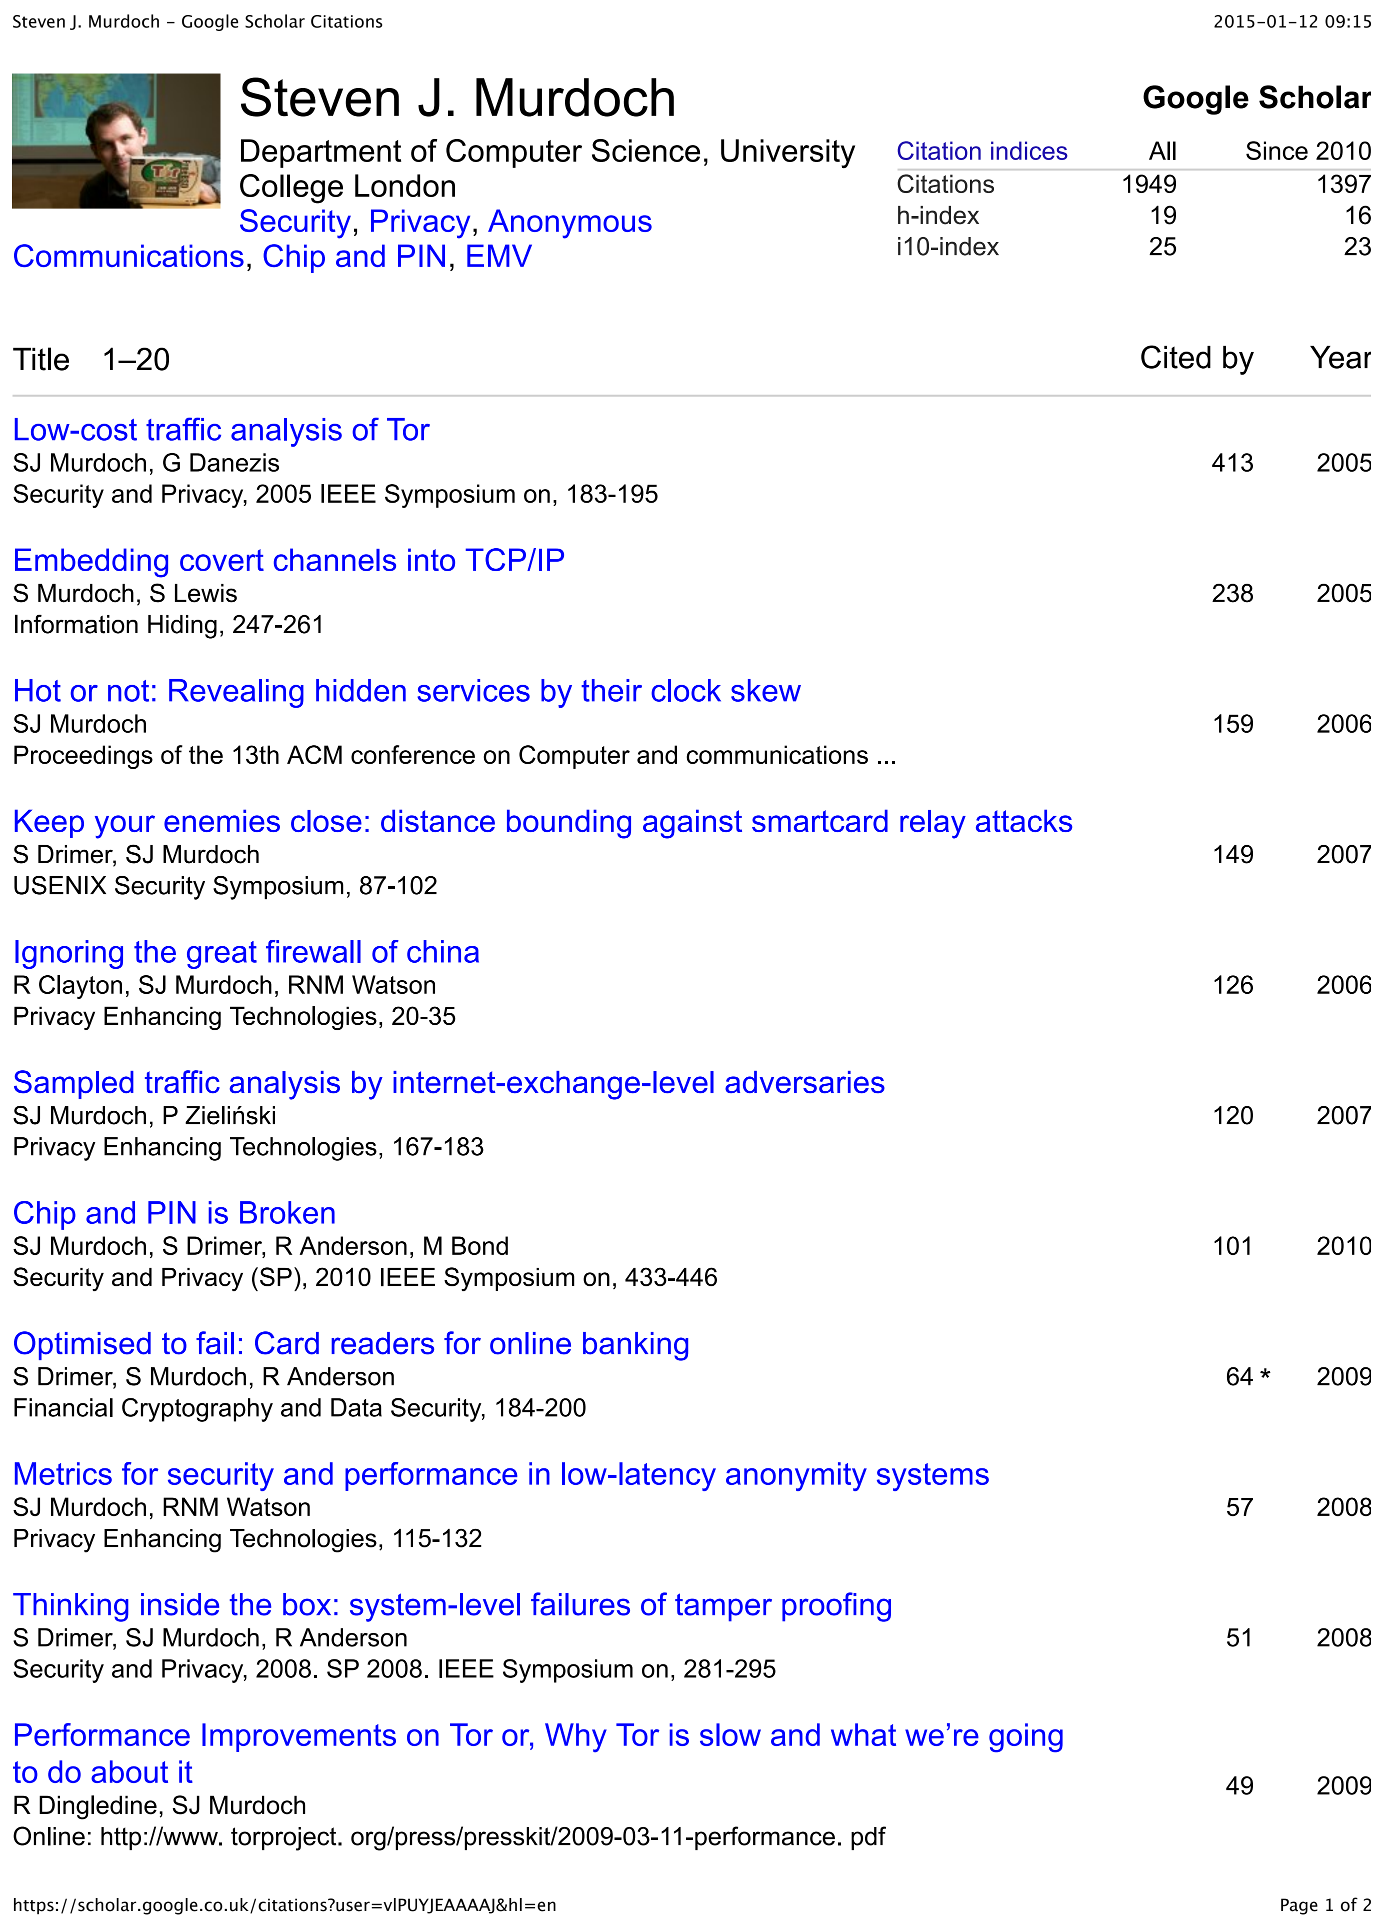

Steven J. Murdoch

Google Scholar

Department of Computer Science, University College London

Security, Privacy, Anonymous Communications, Chip and PIN, EMV

Citation indices	All	Since 2010
Citations	1949	1397
h-index	19	16
i10-index	25	23

Title 1–20	Cited by	Year
Low-cost traffic analysis of Tor SJ Murdoch, G Danezis Security and Privacy, 2005 IEEE Symposium on, 183-195	413	2005
Embedding covert channels into TCP/IP S Murdoch, S Lewis Information Hiding, 247-261	238	2005
Hot or not: Revealing hidden services by their clock skew SJ Murdoch Proceedings of the 13th ACM conference on Computer and communications ...	159	2006
Keep your enemies close: distance bounding against smartcard relay attacks S Drimer, SJ Murdoch USENIX Security Symposium, 87-102	149	2007
Ignoring the great firewall of china R Clayton, SJ Murdoch, RNM Watson Privacy Enhancing Technologies, 20-35	126	2006
Sampled traffic analysis by internet-exchange-level adversaries SJ Murdoch, P Zieliński Privacy Enhancing Technologies, 167-183	120	2007
Chip and PIN is Broken SJ Murdoch, S Drimer, R Anderson, M Bond Security and Privacy (SP), 2010 IEEE Symposium on, 433-446	101	2010
Optimised to fail: Card readers for online banking S Drimer, S Murdoch, R Anderson Financial Cryptography and Data Security, 184-200	64 *	2009
Metrics for security and performance in low-latency anonymity systems SJ Murdoch, RNM Watson Privacy Enhancing Technologies, 115-132	57	2008
Thinking inside the box: system-level failures of tamper proofing S Drimer, SJ Murdoch, R Anderson Security and Privacy, 2008. SP 2008. IEEE Symposium on, 281-295	51	2008
Performance Improvements on Tor or, Why Tor is slow and what we're going to do about it R Dingledine, SJ Murdoch Online: http://www. torproject. org/press/presskit/2009-03-11-performance. pdf	49	2009

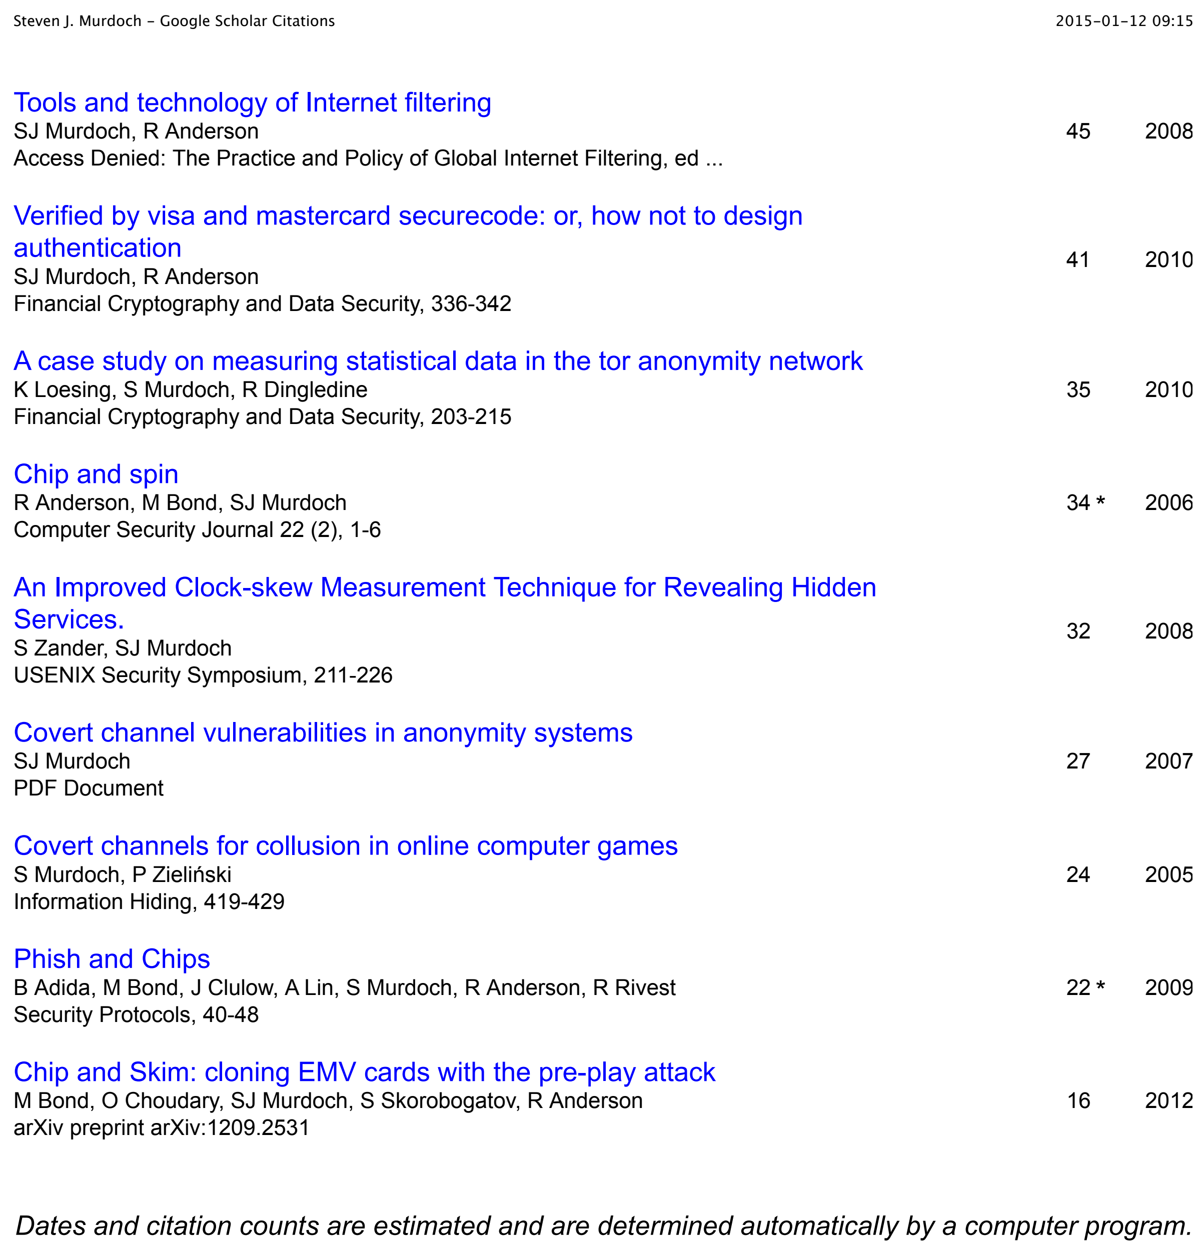

Title	Cited by	Year
Tools and technology of Internet filtering SJ Murdoch, R Anderson Access Denied: The Practice and Policy of Global Internet Filtering, ed ...	45	2008
Verified by visa and mastercard securecode: or, how not to design authentication SJ Murdoch, R Anderson Financial Cryptography and Data Security, 336-342	41	2010
A case study on measuring statistical data in the tor anonymity network K Loesing, S Murdoch, R Dingledine Financial Cryptography and Data Security, 203-215	35	2010
Chip and spin R Anderson, M Bond, SJ Murdoch Computer Security Journal 22 (2), 1-6	34 *	2006
An Improved Clock-skew Measurement Technique for Revealing Hidden Services. S Zander, SJ Murdoch USENIX Security Symposium, 211-226	32	2008
Covert channel vulnerabilities in anonymity systems SJ Murdoch PDF Document	27	2007
Covert channels for collusion in online computer games S Murdoch, P Zieliński Information Hiding, 419-429	24	2005
Phish and Chips B Adida, M Bond, J Clulow, A Lin, S Murdoch, R Anderson, R Rivest Security Protocols, 40-48	22 *	2009
Chip and Skim: cloning EMV cards with the pre-play attack M Bond, O Choudary, SJ Murdoch, S Skorobogatov, R Anderson arXiv preprint arXiv:1209.2531	16	2012

Dates and citation counts are estimated and are determined automatically by a computer program.

Peer review

- An expert in the field reads the paper
- Time consuming, subjective and expensive
- Probably best way to achieve goals
- Used by Research Excellence Framework

Understanding a paper

- Have conclusions been properly drawn?
- Has data been collected and processed in an appropriate way?
- Were experiments done properly (if appropriate)?
- What assumptions were made?
- What other papers should you read to learn more?

How to Read a Paper

S. Keshav
David R. Cheriton School of Computer Science, University of Waterloo
Waterloo, ON, Canada
keshav@uwaterloo.ca

ABSTRACT

Researchers spend a great deal of time reading research papers. However, this skill is rarely taught, leading to much wasted effort. This article outlines a practical and efficient *three-pass method* for reading research papers. I also describe how to use this method to do a literature survey.

Categories and Subject Descriptors: A.1 [Introductory and Survey]

General Terms: Documentation.

Keywords: Paper, Reading, Hints.

1. INTRODUCTION

Researchers must read papers for several reasons: to review them for a conference or a class, to keep current in their field, or for a literature survey of a new field. A typical researcher will likely spend hundreds of hours every year reading papers.

Learning to efficiently read a paper is a critical but rarely taught skill. Beginning graduate students, therefore, must learn on their own using trial and error. Students waste much effort in the process and are frequently driven to frustration.

For many years I have used a simple approach to efficiently read papers. This paper describes the 'three-pass' approach

4. Glance over the references, mentally ticking off the ones you've already read

At the end of the first pass, you should be able to answer the *five Cs*:

1. *Category*: What type of paper is this? A measurement paper? An analysis of an existing system? A description of a research prototype?

2. *Context*: Which other papers is it related to? Which theoretical bases were used to analyze the problem?

3. *Correctness*: Do the assumptions appear to be valid?

4. *Contributions*: What are the paper's main contributions?

5. *Clarity*: Is the paper well written?

Using this information, you may choose not to read further. This could be because the paper doesn't interest you, or you don't know enough about the area to understand the paper, or that the authors make invalid assumptions. The first pass is adequate for papers that aren't in your research area, but may someday prove relevant.

Incidentally, when you write a paper, you can expect most

2. THE THREE-PASS APPROACH

The key idea is that you should read the paper in up to three passes, instead of starting at the beginning and plowing your way to the end. Each pass accomplishes specific goals and builds upon the previous pass: The *first* pass gives you a general idea about the paper. The *second* pass lets you grasp the paper's content, but not its details. The *third* pass helps you understand the paper in depth.

2.1 The first pass

The first pass is a quick scan to get a bird's-eye view of the paper. You can also decide whether you need to do any more passes. This pass should take about five to ten minutes and consists of the following steps:

1. Carefully read the title, abstract, and introduction
2. Read the section and sub-section headings, but ignore everything else
3. Read the conclusions

care to choose coh
to write concise anc
cannot understand
likely be rejected; i
lights of the paper
never be read.

2.2 The secon

In the second pas
ignore details such
points, or to make

1. Look carefully
trations in the
Are the axes p
error bars, sc
nificant? Cor
rushed, shodd
2. Remember to
ther reading (
the backgroun

vey.

PPROACH

read the paper in up to
the beginning and plow-
ass accomplishes specific
us pass: The *first* pass
paper. The *second* pass
, but not its details. The
ne paper in depth.

get a bird's-eye view of
ether you need to do any
about five to ten minutes

ract, and introduction

tion headings, but ignore

reviewers (and readers) to make only one pass over it. Take care to choose coherent section and sub-section titles and to write concise and comprehensive abstracts. If a reviewer cannot understand the gist after one pass, the paper will likely be rejected; if a reader cannot understand the highlights of the paper after five minutes, the paper will likely never be read.

2.2 The second pass

In the second pass, read the paper with greater care, but ignore details such as proofs. It helps to jot down the key points, or to make comments in the margins, as you read.

1. Look carefully at the figures, diagrams and other illustrations in the paper. Pay special attention to graphs. Are the axes properly labeled? Are results shown with error bars, so that conclusions are statistically significant? Common mistakes like these will separate rushed, shoddy work from the truly excellent.

2. Remember to mark relevant unread references for further reading (this is a good way to learn more about the background of the paper).

2.3 The third pass

To fully understand a paper, particularly if you are reviewer, requires a third pass. The key to the third pass is to attempt to *virtually re-implement* the paper: that is, making the same assumptions as the authors, re-create the work. By comparing this re-creation with the actual paper, you can easily identify not only a paper's innovations, but also its hidden failings and assumptions.

This pass requires great attention to detail. You should identify and challenge every assumption in every statement. Moreover, you should think about how you yourself would present a particular idea. This comparison of the actual with the virtual lends a sharp insight into the proof and presentation techniques in the paper and you can very likely add this to your repertoire of tools. During this pass, you should also jot down ideas for future work.

This pass can take about four or five hours for beginners, and about an hour for an experienced reader. At the end of this pass, you should be able to reconstruct the entire structure of the paper from memory, as well as be able to identify its strong and weak points. In particular, you should be able to pinpoint implicit assumptions, missing citations to relevant work, and potential issues with experimental or analytical techniques.

Module Assessment

- You will choose a set of three papers
  - One for presentation in class
  - Two for review
- Choices are constrained for fairness and to give a diverse range of topics
- To maintain fairness, marks will be calibrated depending on:
  - Whether it is an early or a late (in the course) presentation/review
  - The difficulty of the paper

Presentations

- Presentation slides to be submitted on Moodle by **10am on day of presentation**, in PDF format
- As a minimum, you must present most important parts, principal strengths and weaknesses, ethical concerns (if any), and use (if appropriate) of the scientific method
- Maximum time: 15 minutes (will be enforced)

Presentations

- Critically engage with the paper you are presenting – Do not just summarise it
- Assume audience has taken Introduction to Cryptography and Computer Security I
- Try to present something new/interesting
- Make presentation easy to follow and engaging
- Practice alone, then practice in front of friends

Discussions

- After each presentation the class will be invited to ask the speaker questions and engage in a discussion, particularly those who reviewed the paper
- To be able to properly discuss the paper, read the abstract and conclusion of the papers being presented and skim other parts
- Say what was good about the presentations and what could be improved

Paper review

- One page (form and instructions will be on Moodle)
  - Summary of the problem and description of the contribution.
  - The best about the paper for instance new ideas, proofs, simplifications, formalizations, implementation, performance improvement, new insight, expected impact of paper on society, etc.
  - Weaknesses of the paper for instance lack of originality, small increment over previous work, unsubstantiated claims, bad presentation, insufficient discussion of relation with prior work, etc.
  - Put the work in context of the field and discuss its contribution
  - Grade (should it be accepted for publication)
- Due at 10am on day of presentation (same as slides)

Assignment of papers

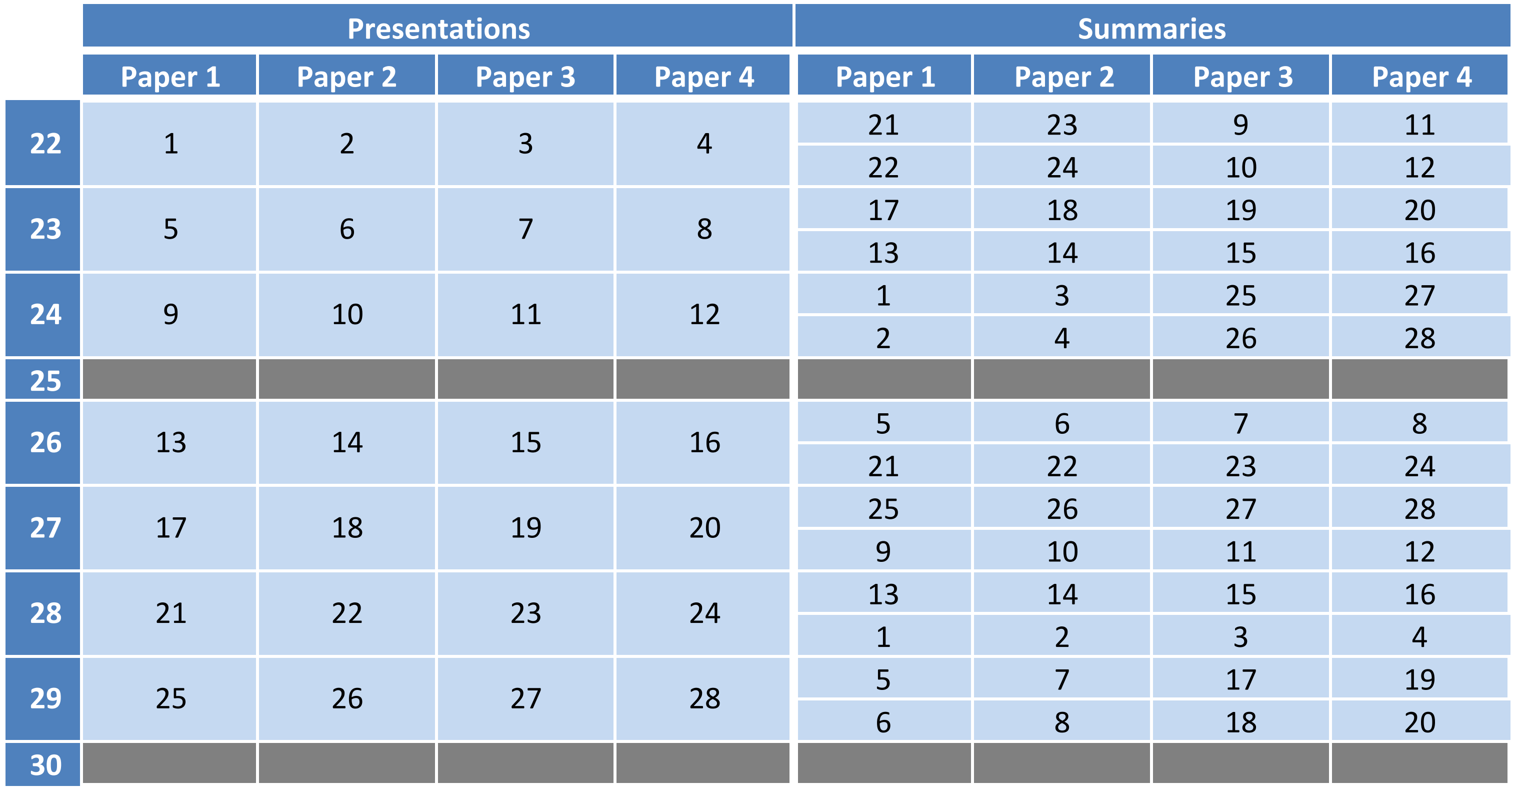

	Presentations				Summaries			
	Paper 1	Paper 2	Paper 3	Paper 4	Paper 1	Paper 2	Paper 3	Paper 4
22	1	2	3	4	21	23	9	11
					22	24	10	12
23	5	6	7	8	17	18	19	20
					13	14	15	16
24	9	10	11	12	1	3	25	27
					2	4	26	28
25								
26	13	14	15	16	5	6	7	8
					21	22	23	24
27	17	18	19	20	25	26	27	28
					9	10	11	12
28	21	22	23	24	13	14	15	16
					1	2	3	4
29	25	26	27	28	5	7	17	19
					6	8	18	20
30								

- You must do one presentation and two paper summaries
- All must be on different topics
- Choose a number and select from questionnaire on Moodle, available after the lecture and to be completed by **10am on Tuesday 17 January**
- The order in students submit the questionnaire is not significant, so there is no rush to complete

Assignment of papers

	Presentations				Summaries			
	Paper 1	Paper 2	Paper 3	Paper 4	Paper 1	Paper 2	Paper 3	Paper 4
22	1	2	3	4	21	23	9	11
					22	24	10	12
23	5	6	7	8	17	18	19	20
					13	14	15	16
24	9	10	11	12	1	3	25	27
					2	4	26	28
25								
26	13	14	15	16	5	6	7	8
					21	22	23	24
27	17	18	19	20	25	26	27	28
					9	10	11	12
28	21	22	23	24	13	14	15	16
					1	2	3	4
29	25	26	27	28	5	7	17	19
					6	8	18	20
30								

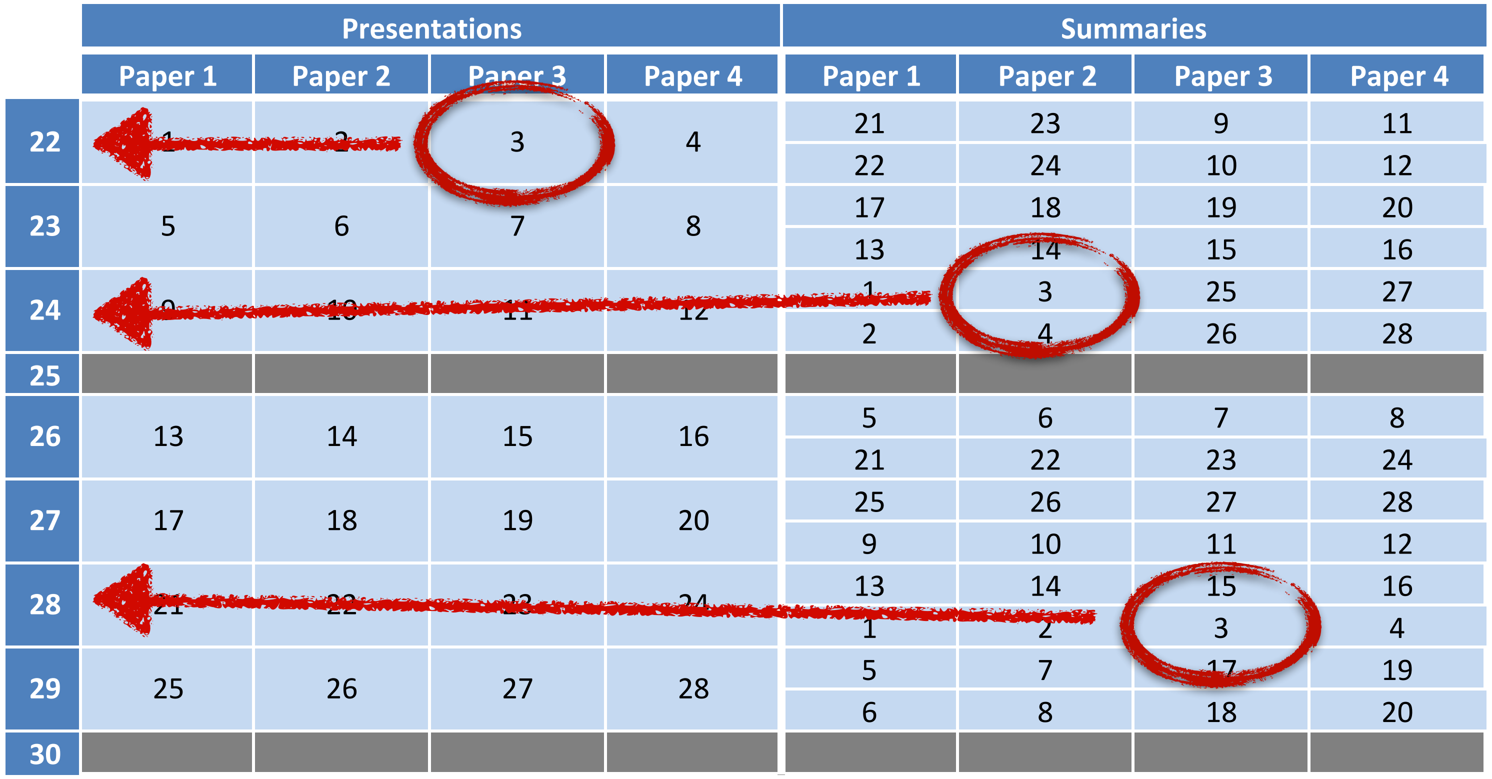

	Presentations				Summaries			
	Paper 1	Paper 2	Paper 3	Paper 4	Paper 1	Paper 2	Paper 3	Paper 4
22	1	2	3	4	21	23	9	11
					22	24	10	12
23	5	6	7	8	17	18	19	20
					13	14	15	16
24	9	10	11	12	1	3	25	27
					2	4	26	28
25								
26	13	14	15	16	5	6	7	8
					21	22	23	24
27	17	18	19	20	25	26	27	28
					9	10	11	12
28	21	22	23	24	13	14	15	16
					1	2	3	4
29	25	26	27	28	5	7	17	19
					6	8	18	20
30								

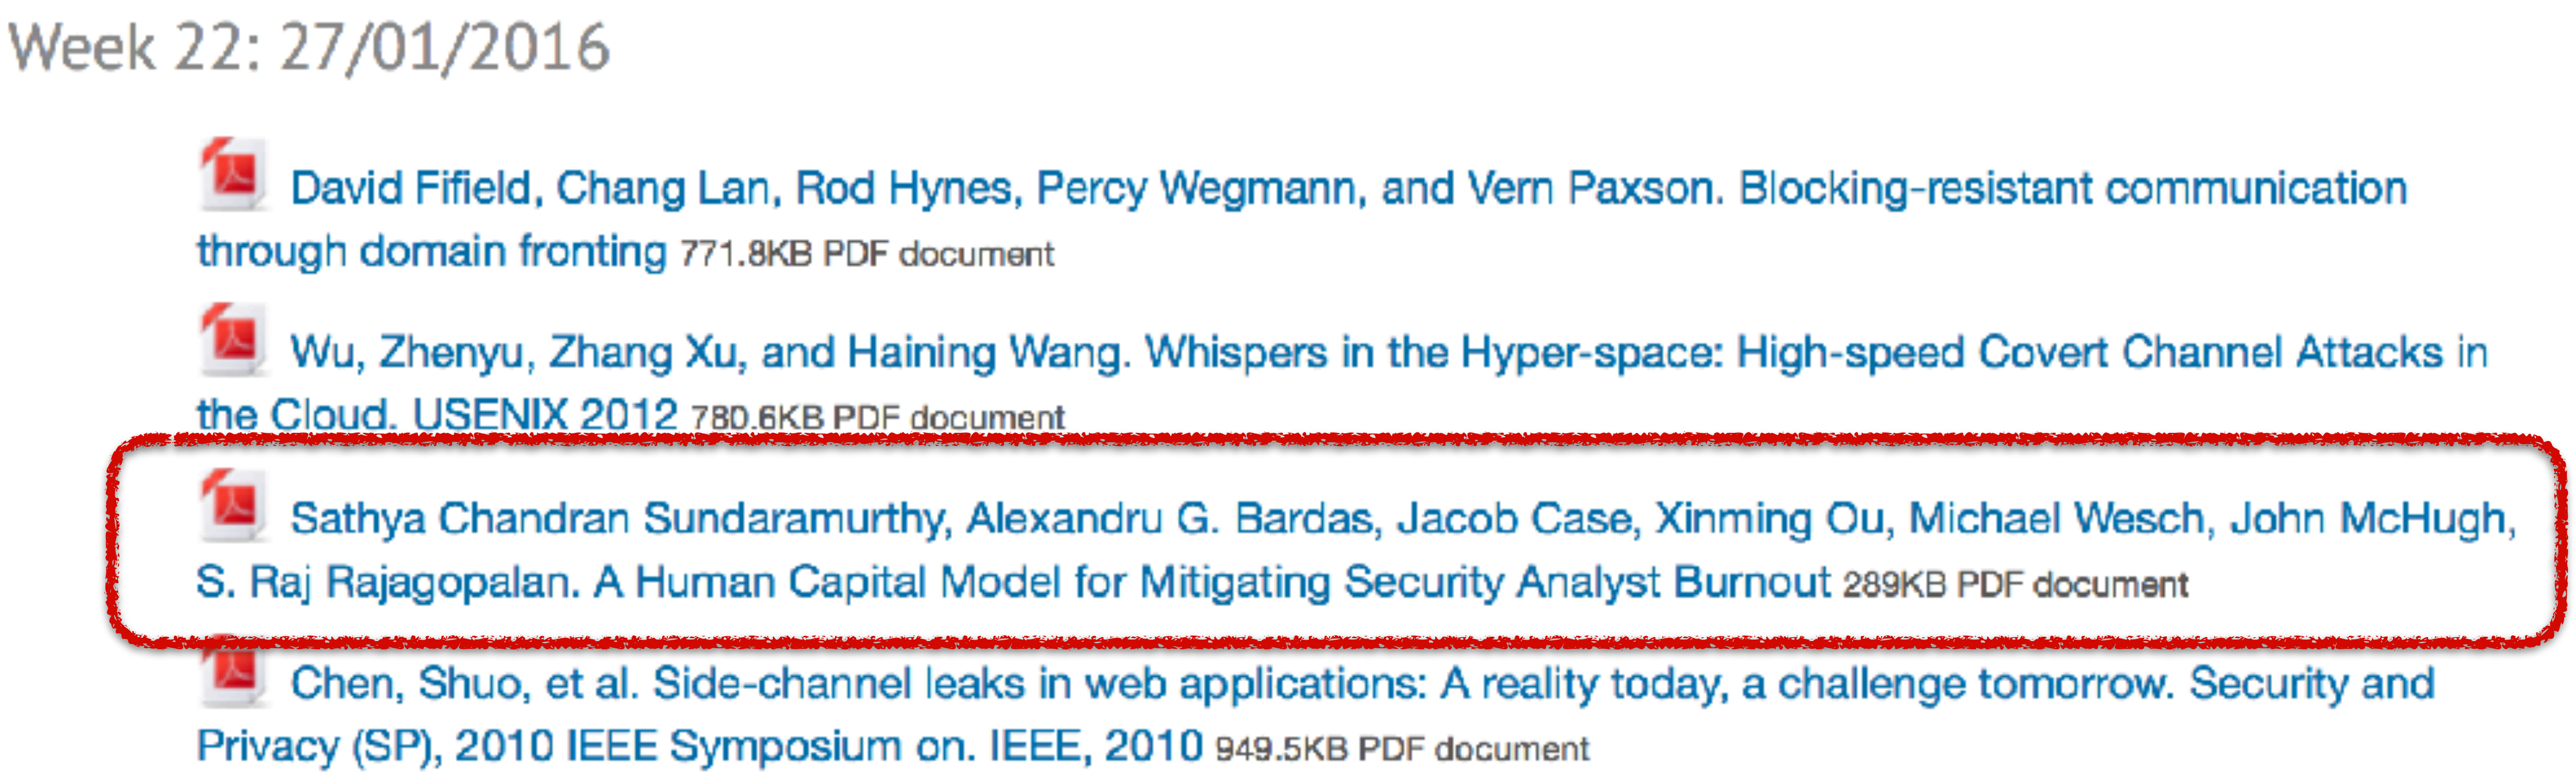

Week 22: 27/01/2016

David Fifield, Chang Lan, Rod Hynes, Percy Wegmann, and Vern Paxson. Blocking-resistant communication through domain fronting 771.8KB PDF document

Wu, Zhenyu, Zhang Xu, and Haining Wang. Whispers in the Hyper-space: High-speed Covert Channel Attacks in the Cloud. USENIX 2012 780.6KB PDF document

Sathya Chandran Sundaramurthy, Alexandru G. Bardas, Jacob Case, Xinming Ou, Michael Wesch, John McHugh, S. Raj Rajagopalan. A Human Capital Model for Mitigating Security Analyst Burnout 289KB PDF document

Chen, Shuo, et al. Side-channel leaks in web applications: A reality today, a challenge tomorrow. Security and Privacy (SP), 2010 IEEE Symposium on. IEEE, 2010 949.5KB PDF document

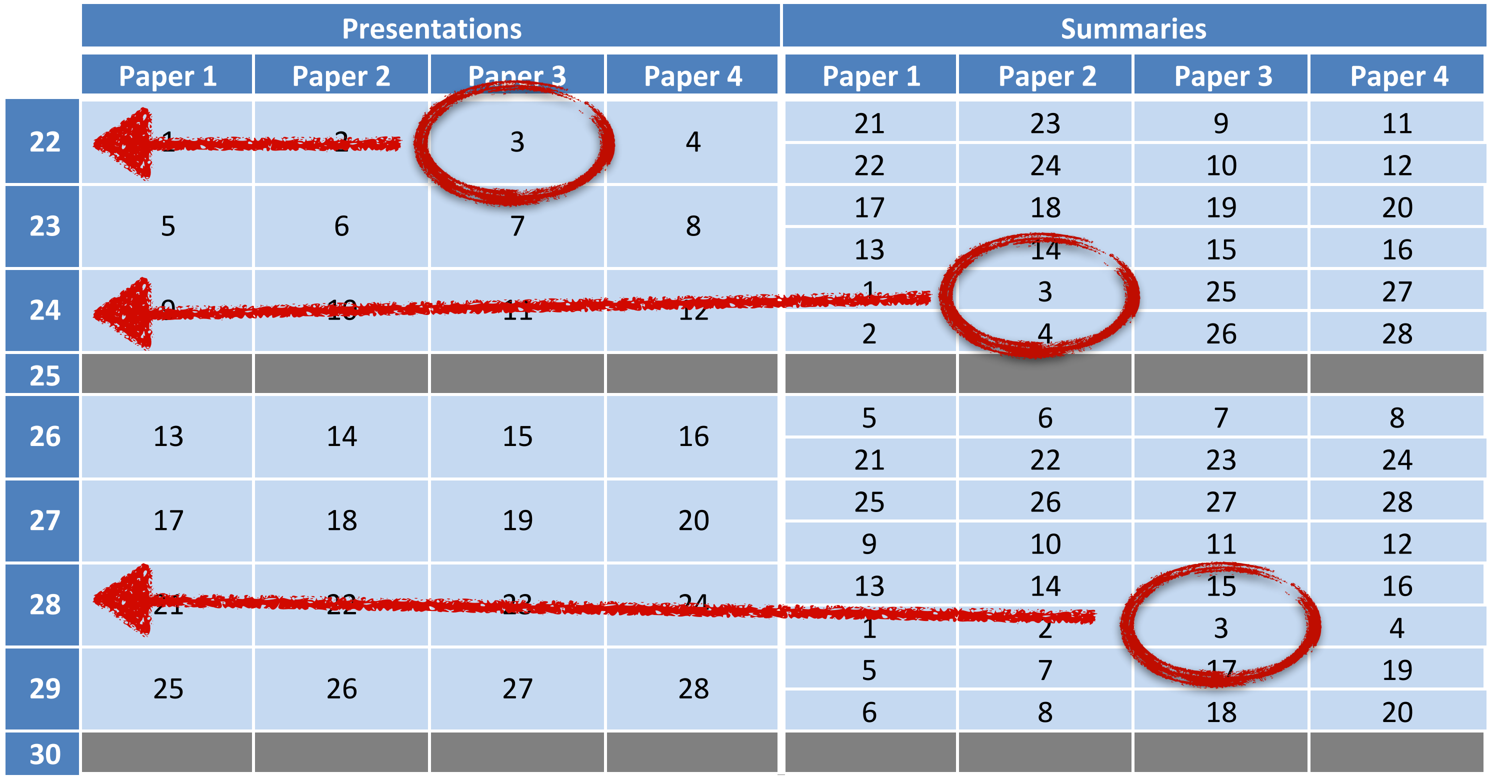

	Presentations				Summaries			
	Paper 1	Paper 2	Paper 3	Paper 4	Paper 1	Paper 2	Paper 3	Paper 4
22	1	2	3	4	21	23	9	11
					22	24	10	12
23	5	6	7	8	17	18	19	20
					13	14	15	16
24	9	10	11	12	1	3	25	27
					2	4	26	28
25								
26	13	14	15	16	5	6	7	8
					21	22	23	24
27	17	18	19	20	25	26	27	28
					9	10	11	12
28	21	22	23	24	13	14	15	16
					1	2	3	4
29	25	26	27	28	5	7	17	19
					6	8	18	20
30								

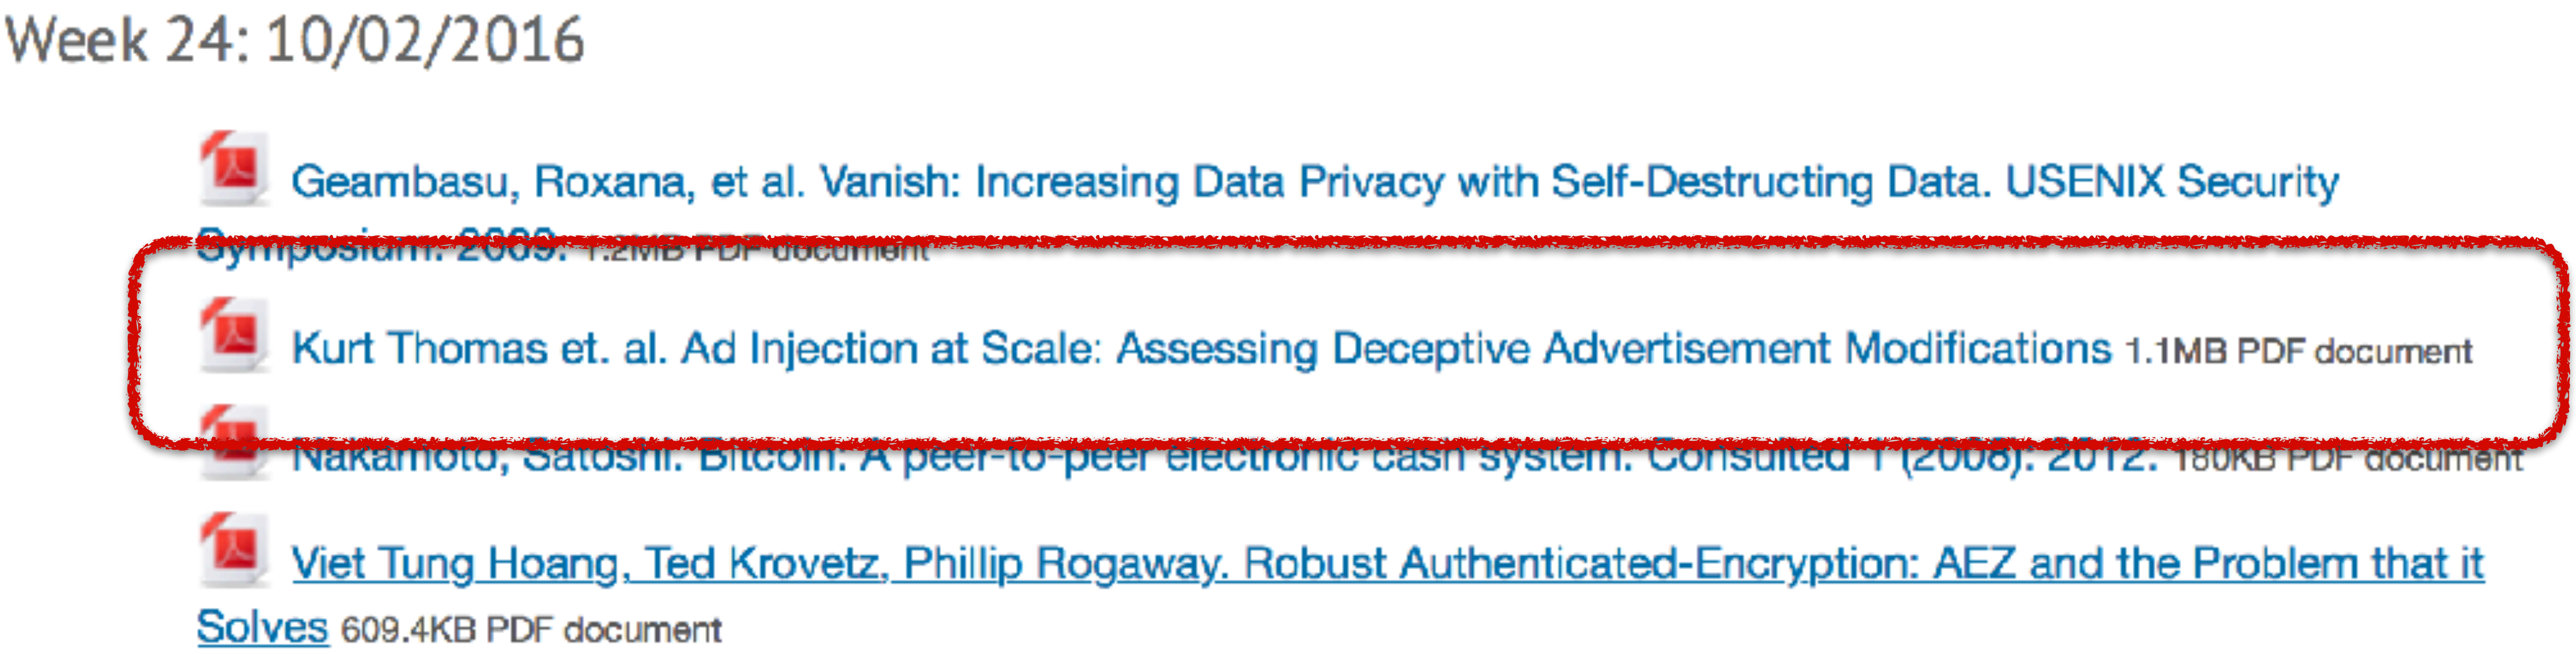

Week 24: 10/02/2016

- Geambasu, Roxana, et al. Vanish: Increasing Data Privacy with Self-Destructing Data. USENIX Security Symposium. 2009. 1.2MB PDF document
- Kurt Thomas et. al. Ad Injection at Scale: Assessing Deceptive Advertisement Modifications 1.1MB PDF document
- Nakamoto, Satoshi. Bitcoin: A peer-to-peer electronic cash system. Consulted 1 (2008): 2012. 180KB PDF document
- Viet Tung Hoang, Ted Krovetz, Phillip Rogaway. Robust Authenticated-Encryption: AEZ and the Problem that it Solves 609.4KB PDF document

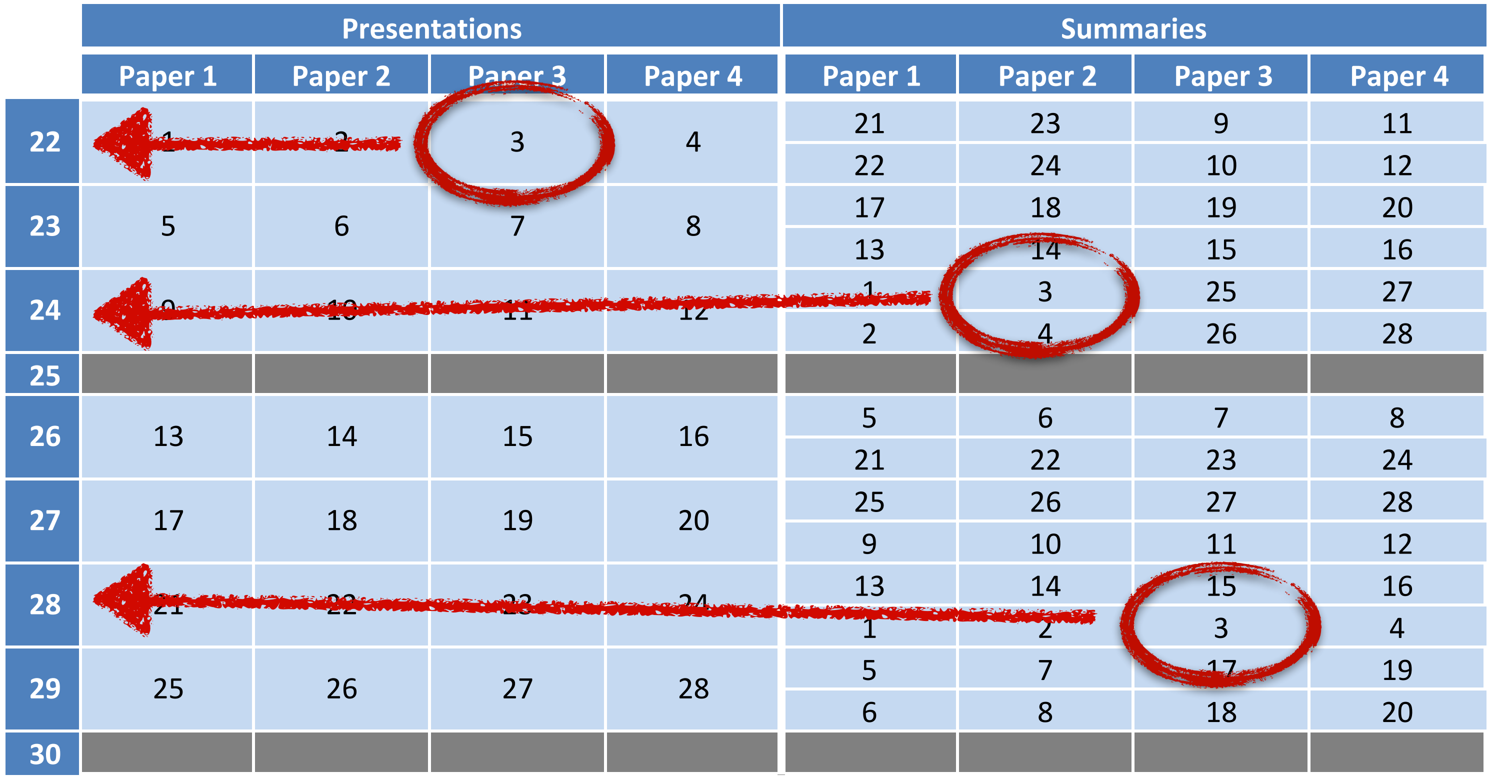

	Presentations				Summaries			
	Paper 1	Paper 2	Paper 3	Paper 4	Paper 1	Paper 2	Paper 3	Paper 4
22	1	2	3	4	21	23	9	11
					22	24	10	12
23	5	6	7	8	17	18	19	20
					13	14	15	16
24	9	10	11	12	1	3	25	27
					2	4	26	28
25								
26	13	14	15	16	5	6	7	8
					21	22	23	24
27	17	18	19	20	25	26	27	28
					9	10	11	12
28	21	22	23	24	13	14	15	16
					1	2	3	4
29	25	26	27	28	5	7	17	19
					6	8	18	20
30								

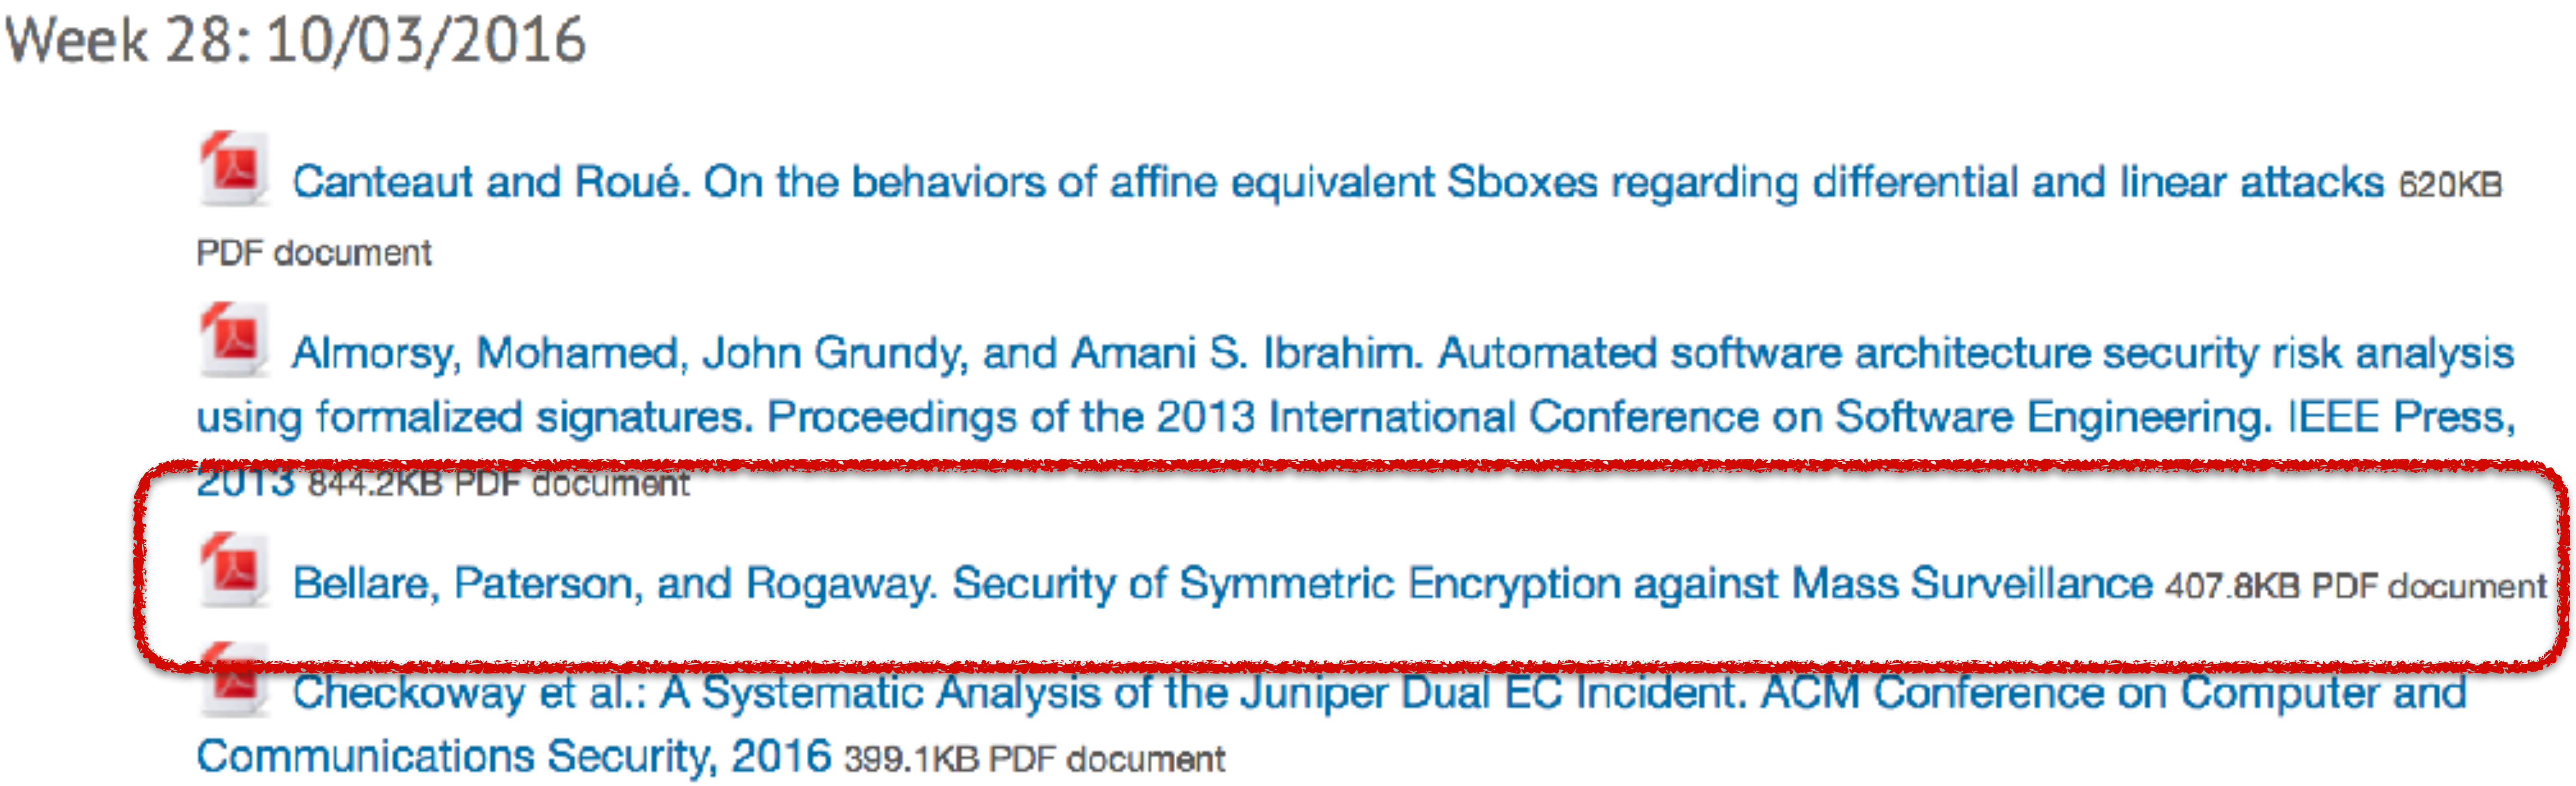

Week 28: 10/03/2016

- Canteaut and Roué. On the behaviors of affine equivalent Sboxes regarding differential and linear attacks 620KB PDF document
- Almorsy, Mohamed, John Grundy, and Amani S. Ibrahim. Automated software architecture security risk analysis using formalized signatures. Proceedings of the 2013 International Conference on Software Engineering. IEEE Press, 2013 844.2KB PDF document
- Bellare, Paterson, and Rogaway. Security of Symmetric Encryption against Mass Surveillance 407.8KB PDF document
- Checkoway et al.: A Systematic Analysis of the Juniper Dual EC Incident. ACM Conference on Computer and Communications Security, 2016 399.1KB PDF document

Literature survey

- The aim of a literature review (sometimes called a literature survey) is to demonstrate to the reader that you have read and understood the main published work concerning a particular topic, and can summarise it, and objectively and critically review it.

Literature survey

- Due Wednesday April 26th 2017 at 5pm (but remember exam preparation)
- Can be about topic of your MSc Information Security dissertation
    - Cannot be copied into your dissertation, but will be a useful foundation
    - If dissertation is done by a pair, so can your survey
    - 20 pages (individual) or 35 pages (pair)
- Otherwise can be on topic of one paper presented in course

More on assessment and feedback for this course

- Submit slides and paper summaries by 10am on the day that the paper is to be presented
- Marks and feedback will be sent to student **within 2 weeks** of the submission
- The student work and corresponding feedback will be made available to all class members on Moodle (but not the marks)
- Literature review will be submitted after the end of the course and feedback will be **within 4 weeks of submission**