serene risc Digest Summer 2016

Cutting Edge Research Summaries for Policy-Makers and Practitioners

Can hackers trust their own reputation systems?

It is profitable for online offenders to cooperate with other skilled associates. Cooperation between offenders is different online than offline and trust is difficult to create and maintain in cyberspace. Understanding how trust works in online criminal communities can help with figuring out how offender groups operate together online. Dupont et al. looked at interactions between users on a massive hacking forum with a reputation system where members can give negative or positive ratings to other members. From the discussion on this forum, they extracted and categorized nearly 450 000 ratings and 25 000 random comments. They found that the existence of a reputation system does not mean that it will be used by all. Only a tiny fraction of forum members participated in the feedback system. New and mid-level hackers disproportionately provided positive feedback and a closer look at the comments revealed arbitrarily sarcastic or silly justifications for positive ratings. This imbalances the reputation system and suggests that it might not be an adequate trust-building tool for criminal communities. Given that a small group maintains the trust mechanism of the forum, police interventions should consider focussing resources on these core members.

Are amateur online crowd investigations helpful to police?

The Boston Marathon bombing was followed by an important moment for public participation in online police investigations. Civilians ran terrorism investigations on the online community Reddit alongside the formal police case. Nhan, Huey and Broll studied over 20 000 of the most popular comments in the aftermath of the Boston Marathon bombing on Reddit. Most comments were general and focussed on self-expression, but other posts shared and distributed news, offered help or provided investigative information. The public has potential to work with law enforcement in the aftermath of a tragedy, but the flow of information must go both ways. It can be harmful if civilians bombard the police with misleading or unguided information. It is important for law enforcement to consider the public as a security partner and to focus attention on investigation areas that require the most support.

How do online criminal groups get together?

Criminal groups often start with family or close friends. Are cybercriminals any different? Leukfeldt, Kleemans and Stol looked at how cybercrime groups involved in online banking offenses got started and grew. By analysing police records, they found four ways that cybercrime groups grow. Their growth may be entirely through real world friends and family, by starting with these people as a base group and recruiting additional criminals online, by using online contacts as a base and recruiting local criminals or entirely through online contacts. The vast majority of the networks studied grew solely through their real world contacts but used forums to gain new knowledge, tools or techniques. Offline social connections are still very important for cybercriminal networks, but online forums have made resources for crime very available.

Why are denial-of-service ‘booter services’ so popular?

A ‘booter service’ will overload a computer on the Internet for a fee, slowing it to a crawl or forcing it offline. Conducting denial of service attacks like this is in most cases illegal but there are booter services openly available online. Hutchings and Clayton wanted to learn more about the individuals behind booter service providers. The researchers interviewed or surveyed 13 openly advertised booter service providers. They found a gradual path to becoming a booter service provider. In general, young males susceptible to the peer influence of gamer friends and online communities appear to move from using these services to providing them. The study participants advertised their services on platforms similar to those that introduced them to booting, such as hacker forums and social media. They generally denied that their services were harmful or illegal and were not worried about law enforcement. Although participants considered running a booter service to be ‘easy money,’ they were concerned about companies like PayPal disrupting payments.

Is there an easier way to reverse engineer code?

Reverse engineers use assembly code analysis to determine the function of machine-readable computer programming code. This type of analysis is time-consuming but essential for detecting copied (cloned) code to discover software plagiarism and security issues. Finding ways to automate parts of this process saves time and frustration for reverse engineers. Ding, Fung and Charland identified the challenges for a practical clone search tool and designed a tool that meets these challenges. Their new tool includes adaptive locality sensitive hashing (ALSH) and integrates inexact assembly code and subgraph searching. The result is a tool that is practical, efficient and scalable. They also constructed a labeled one-to-many assembly code clone dataset for benchmarking. Their solution helps reverse engineers analyse code quickly and thoroughly, helping them work more efficiently. The clone search engine and dataset are available as open source tools.

Do cyber risk awareness programs work for children?

Children start to use the Internet at a young age. However, they are unaware of the risks online. Learning interventions seem promising in raising awareness and changing behaviour in children, but the effectiveness of school-based education programs should be assessed. Schilder, Brusselaers and Bogarts studied the effect of a school-based presentation on online risk awareness and behaviour. The researchers gave a 10-minute presentation on online risk awareness to one group of Belgian students and a separate 10-minute presentation on unrelated topics to another group. They tested both groups' online risk behaviour and awareness immediately after the presentation and again after 4 months. The results showed that the short online risk awareness presentations increased risk awareness and safety knowledge but also led to reporting more risky behaviour. This could be because the participants were better equipped to identify risky behaviour. These short in-class presentations could improve online risk awareness in school aged children.

What is cookie stuffing?

A large number of enterprises employ third party marketers (affiliates) to advertise their products. The affiliates make money when a product is sold though their advertising. Cookie stuffing is a type of fraud that occurs when dishonest affiliates secretly insert their cookies into browsers to get a percentage from sales they did not advertise. Chachra, Savage and Voelker investigated cookie-stuffing fraud in affiliate marketing networks. They found that dishonest affiliates had targeted most affiliate programs. It appeared that dishonest affiliates were more likely to target apparel and accessories e-retailers than other merchant groups. Although cookie-stuffing fraud is not common, it is a novel type of fraud and Merchants should watch revenue flow and consider internal affiliate marketing for greater security.

What's wrong with augmented reality?

Augmented reality (AR) technologies add to a live-view of the physical world with interactive virtual objects. Like any new technology, AR presents unique security and privacy risks. McPherson, Jana and Shmatikov analyzed AR apps to assess these risks. They considered threats from five classes of attackers: AR attackers, ad attackers, web attackers, curious AR services, and network attackers. The researchers then classified six new AR security and privacy vulnerabilities and presented engineering guidelines to improve security and functionality for each one. They recommend that AR service providers prioritize the removal of functional and design flaws to improve security. They also suggest the implementation of security and privacy warnings into AR browsers to inform users of the risks involved.

We know what cyber bullying is, don't we?

Canada's current federal, provincial, and municipal definitions and understandings of cyberbullying are at odds with recent research on cyberbullying. Deschamps and McNutt discuss and outline the current problems Canada faces in how it defines and combats cyberbullying at the federal, provincial, and municipal levels of government. Differences in the definition of cyberbullying have led to a mixed bag of policy instruments and applications. There is a need for a common measurement of what classifies a certain behaviour as cyberbullying. Definitions implying that cyberbullying is similar to traditional face-to-face bullying will not provide decision-makers with the information needed to implement effective policies and legislation. Governmental decision-making could benefit from a well-informed, standardized definition of cyberbullying, based on research.

Can Artificial Intelligence help people choose better passwords?

People are not very good at choosing passwords. Systems that evaluate our password choices can help us to select better passwords. Melicher et al. propose using artificial neural networks to guess passwords. Using neural networks to model password choices and measure their strength is not only possible it also offers benefits over current approaches. They are less resource intensive and work better when making high numbers of guesses or with greater password complexity. Neural network based password selection assistance could prove to be valuable in improving security.

The Ecology of Trust among Hackers

Online offenders profit by cooperating with skilled associates to increase their capacity. However, collaboration is at odds with security as betrayal and dishonesty are common among offenders. Cooperation motivation mechanisms that are relied upon by co-offenders in physical spaces to establish trust, such as coercion by threats of violence are less effective or operate differently online. Consequently, it is very difficult to establish and maintain trust online. Hackers are often forced to rely on reputations when trusting others online. Online forums often have participant rating systems to allow users to self regulate and limit deceit. It is important to understand how trust functions in online criminal communities to better understand how offender groups work together.

Dupont et al. studied how and to what extent these reputation systems are used; if they are effective; what determines trustworthiness; and if these systems encourage trust. They looked at interactions between users on the biggest online discussion forum dedicated to hacking. Although illegal behaviour on the forum is explicitly discouraged, it is implicit in the activities conducted and topics discussed, such as expanding, managing and leasing botnets. The forum used a reputation system where members can give negative or positive ratings to other members. The forum follows a hierarchy where members climb ranks as they contribute to the community. They gain greater capacity to rate others as they increase in rank.

Software extracted nearly 450 000 ratings of 29 985 general and botnet hackers by over 9177 peers during the 2 & ¼ year period until December 2011. Each time a forum member rates another they can leave comments. Researchers selected 25 000 of these comments at random and categorized them by theme. This allowed them to look at how trust is expressed in words as well as numbers.

The existence of a reputation system does not mean that it will be used by all. If a large part of the community doesn't contribute, the reputation system can be ineffective and irrelevant, as the reliability of a reputation system depends on the information it presents. Only a tiny fraction of the hacker forum membership (2.4%) participated in the vast majority (75%) of 'trust exchanges.' This imbalance creates large biases in the recorded reputations. The new and mid level hackers, perhaps fearful of retaliation or attracted to the hacker mystique, provided largely positive feedback. This positive outcropping of feedback does not necessarily reflect the true feelings of the group towards a participant, with the majority of the community not exposing their opinion with commentary. An assessment of the content of comments suggests that sarcastic, humorous, or arbitrary justifications for positive rating are more common than technical or business explanations. Even so, the positive nature of ratings still decreased over time suggesting that the reputation system did little to prevent forum trust from decaying.

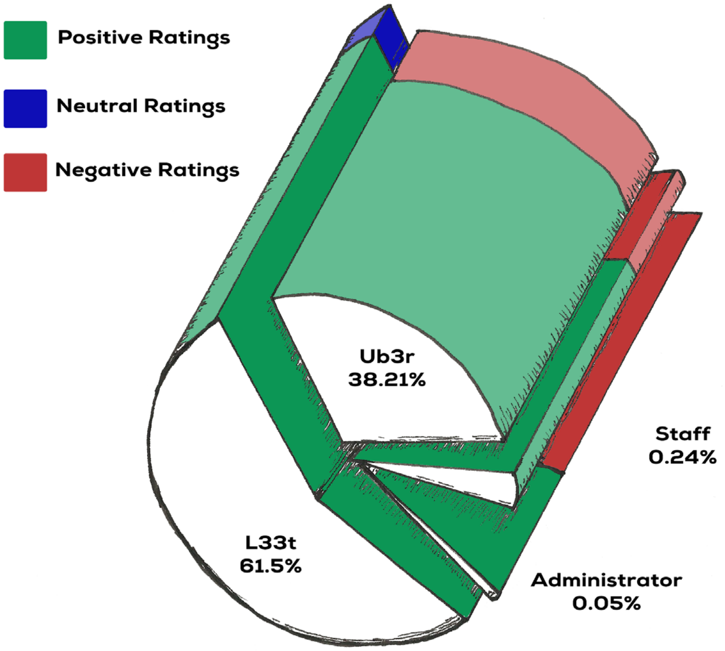

The biases and imbalances identified on this forum suggest that reputation systems may be overrated as trust-building mechanisms for illicit communities.

Police interventions could be improved by focusing resources on the small number of hackers who contribute to reputation ratings, rather than relying on the ratings themselves.

Dupont, D., Côté, A., Savine, C., & Décary-Hétu, D. (2016). "The ecology of trust among hackers." Global Crime, 17(2), 129-151.

Digilantism: An Analysis of Crowdsourcing and the Boston Marathon Bombings

The Boston Marathon bombing marked an important moment for public participation in online police investigations. Civilians ran terrorism investigations on Reddit alongside a formal police case. Reddit is a massive online community where users create and discuss posts on varied subjects and can vote on the popularity of comments. Although the civilian contributions to Reddit failed to identify the perpetrators of the Boston Marathon bombing, they can provide insight on the role that online citizens play in public security.

Nhan, Huey and Broll studied online discussions about the aftermath of the Boston Marathon bombing on Reddit. They explored the public's role as a potential partner to police by highlighting information online helpful to the investigation. During the manhunt for the suspects Reddit users created 20 posts about the bombing. Each post had thousands of posted comments. The researchers read and assigned themes to over 20 000 of the most popular comments in order to pinpoint investigation information.

The majority of comments were very general and focussed on self-expression, but there were also posts sharing and distributing news and other information; offerings of assistance; and investigation information. A small but remarkable number of users contributed to investigation-related comments. Some users provided their own analyses of the crime scene with varying degrees of expertise. Others encouraged reporting pertinent information to the FBI and acted as an extra set of eyes and ears for the police.

It was evident that there is potential for the general public to work alongside law enforcement online when dealing with the aftermath of a tragedy. However, when the flow of information is only from the public to police, it may cause more harm than good. Amateur investigations online may strain police resources by providing a surplus of information. Without proper police guidance, well-intentioned civilians can mislabel innocent actions and individuals as suspicious. Direct interactions and dialogue between police and Reddit users was rare and may suggest a conflict between public security and civil society.

There seems to be a public desire to help the police during their investigations following tragedy. It would be in the best interest of police management to consider online communities as security partners. Law enforcement should consider actively focussing the public's attention on areas of their investigation that require support instead of leaving the public to take matters into their own hands. Collaboration between the police and civilians online may provide unique benefits and curb misidentifications

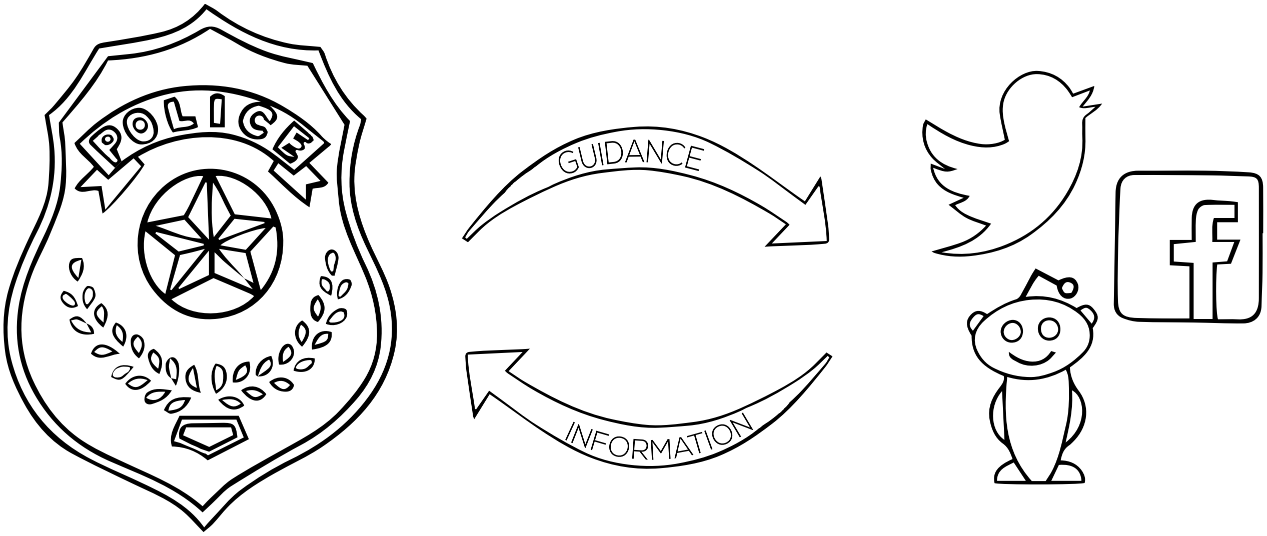

Collaboration between the police and civilians online may provide unique benefits and curb misidentifications.

Nhan, J., Huey, L., & Broll, R. (2015). "Digilantism: An Analysis of Crowdsourcing and the Boston Marathon Bombings." British Journal of Criminology. doi:10.1093/bjc/azv118

Cybercriminal Networks, Social Ties and Online Forums: Social Ties versus Digital Ties within Phishing and Malware Networks

Do cybercriminals meet and collaborate differently? Friends, family and other social contacts are often how criminal groups start. To grow in capacity these groups look outside of their existing social circle. Criminals could only expand their network by meeting contacts in appropriate physical spaces; such as cafés and bars. The Internet provides spaces to connect virtually that are both appropriate and specific to various forms of criminal activity. Criminals can exchange information, learn skills, find support and form relationships in dedicated online forums. However, the existence of Internet relationships does not preclude real world social connections.

Leukfeldt, Kleemans and Stol studied police investigation reports on online banking offenses, such as theft of personal information by phishing and malware to gain a better understanding of the social relationships within cybercrime networks. They looked at the network origins and growth, the relationships between members as well as the role of forums within the network. The researchers studied 18 police files on online banking crimes. Ten were phishing attacks, five were malware attacks and three cases combined both phishing and malware attacks. The cases studied included information from recorded phone calls and Internet traffic, undercover observation, house searches and interviews. Group members were seen as performing one of three roles. 'Core members' initiated and coordinated the attacks: without them, no crime would be planned. 'Enablers' provided the services and technical skills necessary for the crime to happen. Finally, 'money mules' disguised the financial trail left by the crime.

Four types of cybercrime group growth

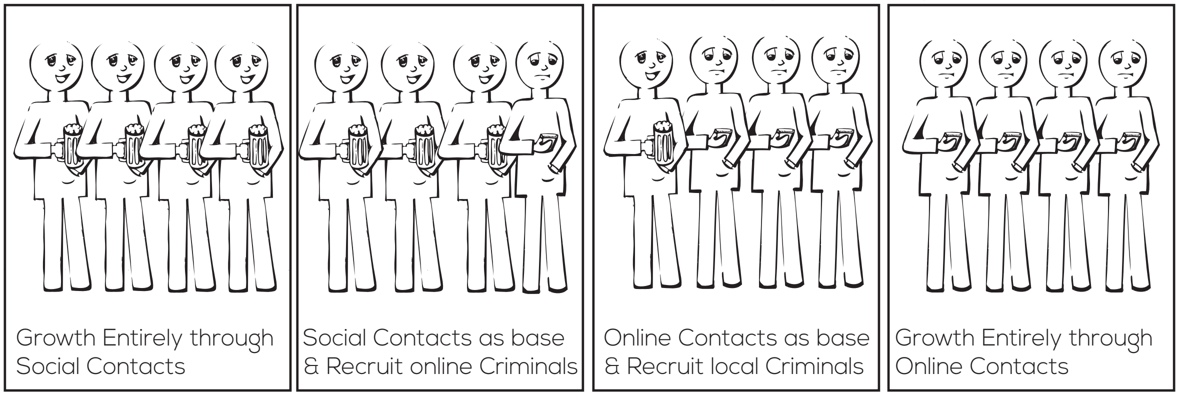
Growth Entirely through Social Contacts
Social Contacts as base & Recruit online Criminals
Online Contacts as base & Recruit local Criminals
Growth Entirely through Online Contacts

70% of the networks grew solely through social contacts, with forums being used only to obtain specific knowledge or tools. Core members recruited enablers from among their family or acquaintances, at sports clubs, schools or on the street. The number and types of enablers that they could use were limited. Enablers recruited were often used for multiple activities out of necessity. 10% of the networks grew from a base of social contacts, but used forums to recruit specialists. They used forums to purchase specific illegal services or products such as mailing lists or malware. One network grew from an online group to include social contacts. Although tools and techniques were acquired and discussed primarily through the Internet, social contacts were used to recruit locals offline as money mules or as service company insiders to overcome security. 10% of networks recruited entirely through forums. They found specialists online, from whom they purchased services and products. They were also able to recruit money mules through spam mailing campaigns.

Social contacts were used to form or grow the majority of the groups studied. However, online forums do impact how groups operate as they seem to make it easier to find accomplices with specific talents and are a key source of criminal knowledge and skills.

Real world social connections are still very important for cybercriminal networks, but online interaction helps with knowledge acquisition.

Leukfeldt, E. R., Kleemans, E. R., & Stol, W. P. (2016). "Cybercriminal networks, social ties and online forums: social ties versus digital ties within phishing and malware networks." British Journal of Criminology. doi:10.1093/bjc/azw009

Exploring the Provision of Online Booter Services

Distributed Denial of Service (DDoS) attacks overload computers until they are unable to operate properly. Online gamers use DDoS attacks to attack their opponents by disrupting their Internet connection or the game itself. This is known as 'booting,' where players 'boot' or remove other players from the game. There is a demand for DDoS attack services and they are provided by websites called 'booter services.' Although there have been studies on the size of the booter market and on individual services, there is a need to explore the backstories of booter service providers.

Hutchings and Clayton wanted to understand more about the individuals behind booter service providers. Their goal was to listen to booter service providers talk about their own lives and activities. The researchers invited booter service providers with publicly accessible websites to take part in a survey. The service providers studied mainly used amplification techniques to overload computers. These techniques exploit a feature of network traffic management to amplify their attack capacity. This allows them to generate enough traffic to overwhelm systems without using a botnet. The researchers communicated with the operators of 63 openly advertised booter services. 13 booter service providers decided to take part in the study. The participants had the opportunity to either fill out an online survey or take part in an interview.

The researchers found that there is a gradual path to becoming a booter services provider. The participants were generally young, male, and influenced by their peers. Many started using booter services as gamers and hackers and became acquainted with the provision of booter services through friends and online communities. Some participants viewed providing booter services as a way to learn new skills. Despite clear indications of illegal use, they tended to deny that they were providing illegal or harmful services, claiming instead that they offered an important service for network testing.

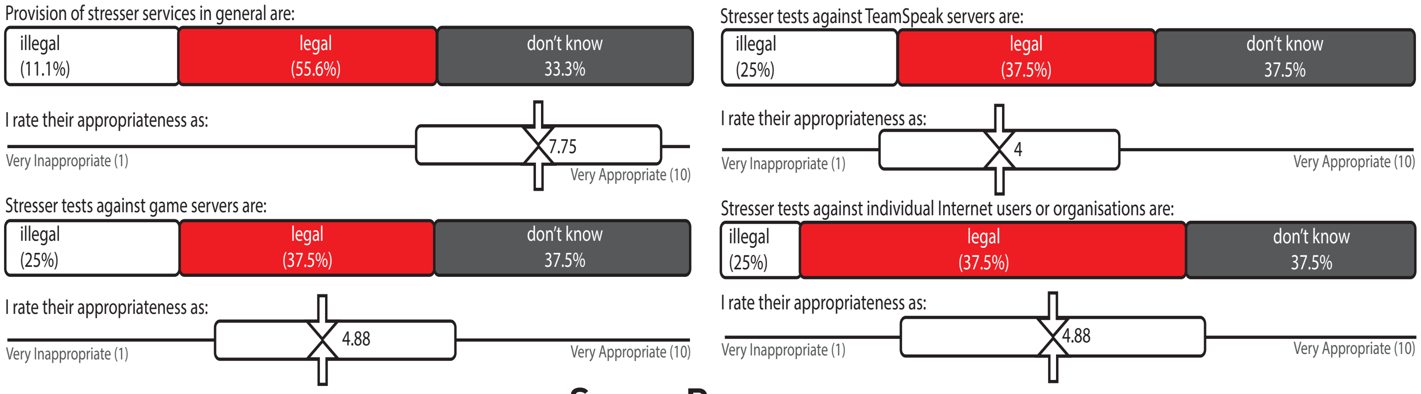

Survey Responses

In general, they advertised their services on platforms similar to those that first introduced them to booting, such as hacker forums and social media. Booter service providers did not seem to be concerned about law enforcement. Their greatest source of frustration was in receiving payment for their services, as Internet transaction companies such as PayPal had disrupted payments to booter service providers from time to time.

Running booter services is currently considered 'easy money,' with little cost in terms of time spent maintaining the sites. The spectre of legitimacy even allows sites to engage large third party companies such as CloudFlare for security to reduce the impact of unfriendly competition. Their main frustration at the moment is receiving payment, since some companies reject their claims of legal operation and deny them service.

Attackers view booter services as profitable, low effort and low risk since they are not often targeted by law enforcement.

Hutchings, A. & Clayton, R. (2016). "Exploring the Provision of Online Booter Services." Deviant Behavior. doi:10.1080/01639625.2016.1169829

Kam1n0: MapReduce-based Assembly Clone Search for Reverse Engineering

Reverse engineering is a time consuming process that involves inspecting computer program code (called assembly code) and determining its function. Analyzing this assembly code is a critical process for detecting software plagiarism and software patent infringements, identifying vulnerabilities and determining the purpose of programs when the source code is unavailable. Reusing program code is common and unregulated. Pieces of code are often reused as they are or modified slightly and then combined to make new programs. Using code in this way generates similar or identical pieces of assembly code or fragments of cloned code. Automating the detection and classification of cloned code is a great help to reverse engineers as it saves them time and limits the amount of code functions they have to decipher manually. Existing approaches to clone searching have focused only on search accuracy. However, in practical applications there are multiple factors that are important for a search tool, such as the speed and responsiveness to the tool.

Ding, Fung and Charland identified a number of challenges for a practical clone search tool. These included the interpretability and usability; efficiency and scalability; ability to update the code library incrementally and the clone search quality provided by the tool. They designed a tool that meets these challenges. The tool makes the search results easier to understand and use by providing subgraph clones as results. This shows small sections of code that are similar to or the same as pieces of code with understood functions. This helps reverse engineers analyze assembly code because they can quickly identify known functions and concentrate on the unknown sections.

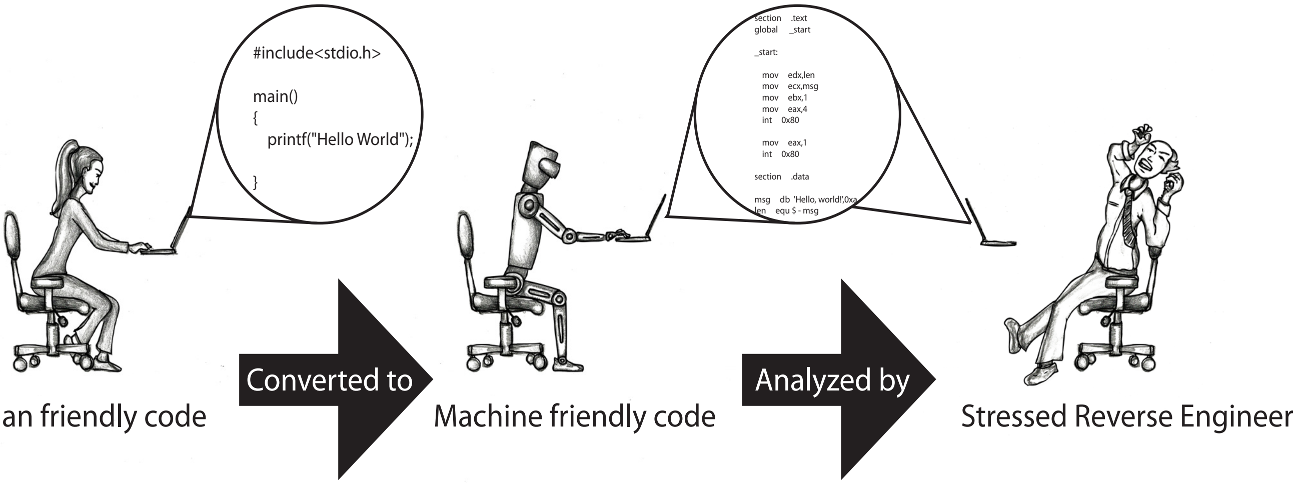

```
#include<stdio.h>

main()
{
   printf("Hello World");

}
```

```
section    .text
global     _start

_start:

   mov    edx,len
   mov    ecx,msg
   mov    ebx,1
   mov    eax,4
   int    0x80

   mov    eax,1
   int    0x80

section    .data

msg    db  'Hello, world!',0xa
len    equ $ - msg
```

The researchers implemented a search method, known as adaptive locality sensitive hashing (ALSH) that is efficient in searching through assembly code. They also applied the first approach that integrates both an inexact assembly code search with a subgraph search to provide high quality search results. They implemented the search tool on the 'Big Data' technologies of MapReduce and the Apache Spark computational framework. To test the system they constructed a labeled one-to-many assembly code clone dataset to allow for benchmarking. This dataset is also available to the research community.

The researchers have demonstrated a solution to help reverse engineers analyse assembly code. They created a clone search engine that is accurate, practical and scalable. It can help reduce the amount of time and effort required to analyse assembly code and allow the engineers to focus on deciphering new elements of code. They have made the clone search and the benchmarking dataset available as open source tools.

Open-source, Big Data search tools can save time and increase reverse engineer performance

Ding. S. H. H., Fung, B. C. M., & Charland, P. (2016). "Kam1n0: MapReduce-based Assembly Clone Search for Reverse Engineering." Presented at KDD 2016.

The Effectiveness of an Intervention to Promote Awareness and Reduce Online Risk Behavior in Early Adolescence

In several European countries the average age at which children start to use the Internet is currently 7-8 years of age; and that age is decreasing. Unfortunately, young children are not aware of the risks they face online. These risks include cyberbullying; the disclosure of personally identifying information; unadulterated exposure to porn, violence, or racism; the abuse of personal information; and spam. The need for online safety is well understood by parents and teachers. This is due to awareness campaigns promoted by local governments and the European Union. Studies have shown that interventions are helpful in raising awareness among youngsters and positively affecting their behaviour. However, there has been a lack of research into the effectiveness of education programs in schools.

This work examines whether school-based interventions are effective in raising online risk awareness and influencing risky behaviour. Schilder, Brusselaers, and Bogaerts conducted a study to measure online risk behaviour and another on online risk awareness. The online risk behaviour study was to know how children act online and whether their behaviour is considered as risky. Some examples of risky behaviours include: meeting strangers; sharing private sensitive information on public profiles; having email addresses or using the Internet without their parent's knowledge; falsifying their identify, or hiding age or gender. The online risk awareness study examined the understanding of the children regarding their safety online. This includes whether the children are aware of the danger of talking with strangers, opening an unknown e-mail attachment, using public social network sites and other risky activities.

The researchers presented awareness materials to students and tested the effectiveness of this approach. They divided 812 Belgians aged between 9 and 11 into two groups. One group was provided with a 10-minute online risk awareness presentation. The other group was provided a 10-minute class on topics not related to online risk. After the presentations, they were surveyed to measure both 'Online Risk Behaviour' and 'Online Risk Awareness.' After 4 months, they were assessed again in the same way. The results of this study are limited by the data collected from the young students not being validated in any way.

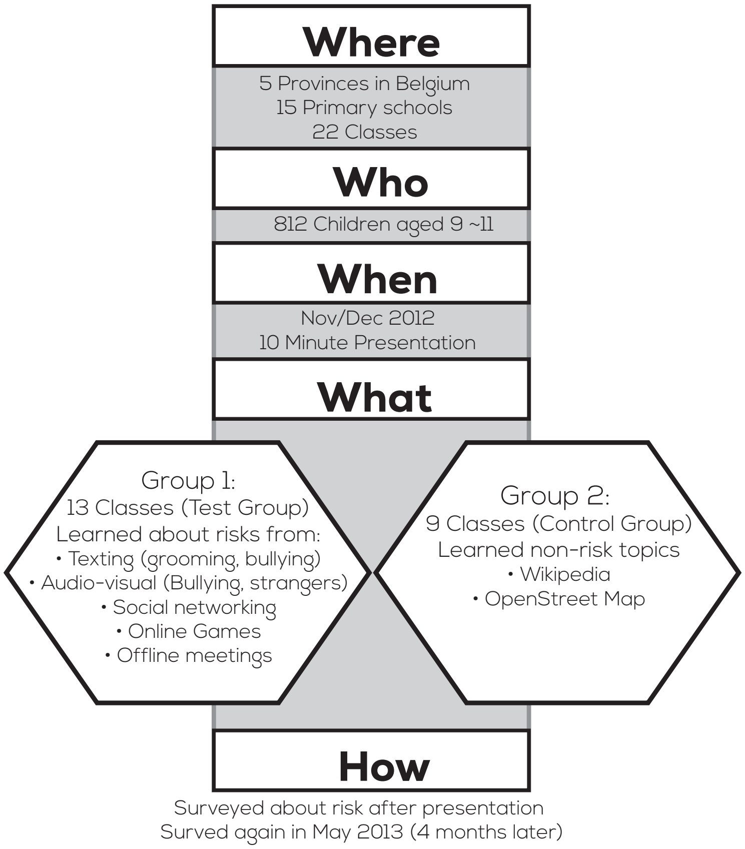

Where

5 Provinces in Belgium
15 Primary schools
22 Classes

Who

812 Children aged 9 ~11

When

Nov/Dec 2012
10 Minute Presentation

What

Group 1:
13 Classes (Test Group)
Learned about risks from:
- Texting (grooming, bullying)
- Audio-visual (Bullying, strangers)
- Social networking
- Online Games
- Offline meetings

Group 2:
9 Classes (Control Group)
Learned non-risk topics
- Wikipedia
- OpenStreet Map

How

Surveyed about risk after presentation
Surved again in May 2013 (4 months later)

The results showed that the short interventions helped to raise the children's online safety knowledge. However, the students reported riskier behaviour after the interventions. This could be because they previously did not know the difference between risky and safe behaviour. After learning about online safety they were more aware of their risky behaviours and thus reported it more. The study also showed that boys and older children were more likely to behave riskily online.

Short educational presentations about online risks influences children's long-term understanding and attitudes towards safe online behaviours.

Schilder, J. D., Brusselaers, M. B. J., & Bogaerts, S. (2015). "The Effectiveness of an Intervention to Promote Awareness and Reduce Online Risk Behavior in Early Adolescence." Journal of Youth and Adolescence, 45(2), 286-300.

Affiliate Crookies: Characterizing Affiliate Marketing Abuse

A large number of enterprises employ third party marketers (affiliates) to advertise their products. The affiliates make money when a product is sold though their advertising. The affiliate depends on web cookies they place on a customer's computer to identify which affiliate has contributed to the sale. A cookie is a piece of data stored in a browser that contains some information. Websites read cookies to recognize users and obtain related information. It is possible for dishonest affiliates to insert their cookies onto computers without the user being aware. The dishonest affiliates will then receive a commission from purchases that have nothing to do with advertising. This is known as 'cookie stuffing' fraud. E-retail businesses lose money due to paying commissions to services that have not increased their revenue.

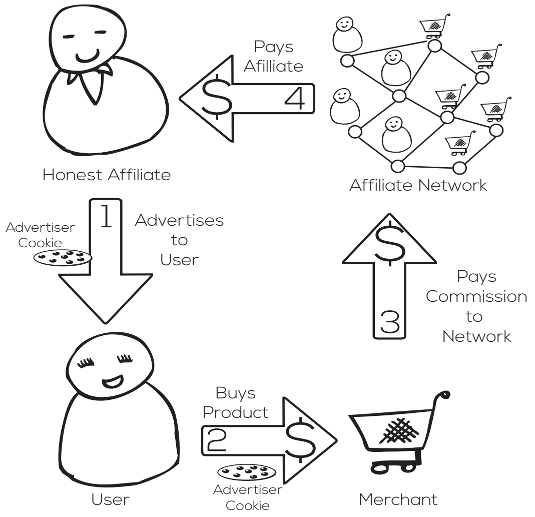

Normal Affiliate Marketing Process

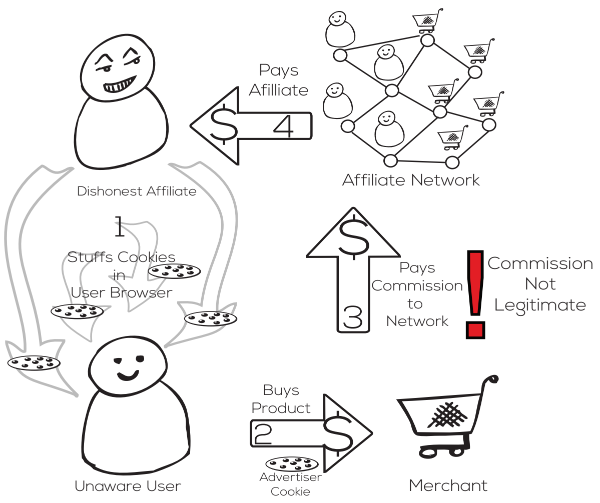

Cookie Stuffing in Affiliate Marketing Process

Chachra, Savage and Voelker investigated cookie-stuffing fraud in affiliate marketing networks. They surveyed affiliate cookies at a large scale using a browser-based tool to identify cookie stuffing.

The researchers chose six large affiliate programs as data domain. An affiliate program is a platform that builds connection between affiliates and organizations. After a purchase is made, the affiliate program extracts cookies from the user's browser and identifies the related affiliates. Researchers employed 74 volunteers to install software to collect affiliate cookies. They classified each of the cookies by affiliate program. Then they used a program that collects web pages to search for stuffed cookies in 475 000 domains.

They found that dishonest affiliates had targeted most affiliate programs. CJ Affiliate and Rakuten LinkShare were the two most severely affected. It appeared that dishonest affiliates were more likely to target apparel and accessories e-retailers than other merchants. The large scale study collected 61 normal affiliate cookies and 0 stuffed cookies during a two-month period. Half of these affiliate cookies were from the Amazon Associates Program. A limited amount of cookie stuffing was observed and a small number of affiliates dominated the affiliate market.

Although cookie-stuffing fraud is not common, it is a novel type of fraud. Affiliate programs and their users do not suffer direct losses from this fraud, but the merchant does. Organizations should pay more attention to affiliate marketing security. If an enterprise is planning to use affiliate marketing, an in-house affiliate marketing program is recommended due to the greater transparency of the affiliate activities and revenue flows.

If your enterprise is using affiliate marketing, keep an eye on revenue flows.

Chachra, N., Savage, S., & Voelker, G. M. (2015) "Affiliate Crookies: Characterizing Affiliate Marketing Abuse." Presented at IMC 2015.

No Escape from Reality: Security and Privacy of Augmented Reality Browsers

Augmented reality (AR) technologies supplement a live view of the real-world with interactive virtual objects. AR applications work by first sensing input, transforming sensed objects, and then displaying transformed objects to the user. A mobile AR browser can combine a live camera feed and GPS data from a smartphone to overlay virtual objects on the camera view in real-time. There are different AR channels that specify the content and how it is displayed, as well as what happens when certain objects come into view or are clicked by the user. AR channels include those that overlay historical pictures when viewing sightseeing landmarks, or display reviews for restaurants.

McPherson, Jana, and Shmatikov conducted an in-depth analysis of AR apps to assess their security and privacy risks. Using five classes of cyber attackers, they classified six types of vulnerabilities and provided corresponding engineering guidelines for securing AR functions. The researchers analysed the functional requirements of AR and tested for new vulnerabilities in their implementation. They then demonstrated new threats unique to AR browsers. The study looked at five classes of attackers: AR attackers, ad attackers, web attackers, curious AR services, and network attackers. They classified six new security and privacy vulnerabilities unique to AR browsers:

1. Access to device resources from outside the AR browser
2. Content from different sources
3. Automatic actions activated by images
4. Images from the user's phone being processed over another network
5. Overlapping visual layers that could cause some content to be invisible to the user
6. Third-party channels in an AR app

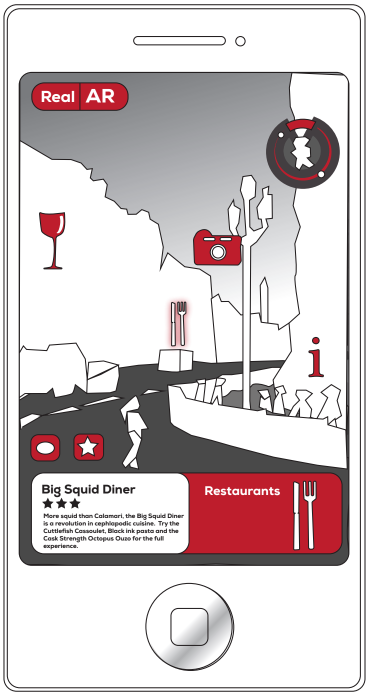
Real AR
Big Squid Diner
★★★
More squid than Calamari, the Big Squid Diner is a revolution in cephlapodic cuisine. Try the Cuttlefish Cassoulet, Black ink pasta and the Cask Strength Octopus Ouzo for the full experience.
Restaurants

Access to device resources from outside the AR browser is possible currently because the AR browser can be accessed by any web content regardless of its origin. This means that attackers can gain access to a user's camera and take pictures of the user and their surroundings. One solution the researchers recommend is that AR browsers apply more restrictions on channels, while giving more control to the user.

Automatic actions activated by images can be a problem in some AR browsers as the browser is constantly analyzing the camera feed. As soon as it recognizes an image associated with a channel, it automatically launches the channel's content without a confirmation prompt. This could lead to fully-automated large scale tracking. For example, registering a channel associated with the image of a vehicle license plate could crowd source tracking of that vehicle. Any time a user scans their surroundings and that license plate comes up in the camera feed, the channel will be launched automatically and the plate's location and surroundings are sent to the channel's owner. AR browsers should inform the user about the origin of the AR content, the possibility of false matches and unexpected content.

AR service providers should prioritize removing functional and design flaws to improve the security of mobile AR browsers. The six categories of AR vulnerabilities and related guidelines for software engineers can help implement better security into existing AR browsers. User experience professionals should also integrate security and privacy warnings into mobile AR browsers so that the users are informed of the risks of their use.

Augmented reality apps present unique security and privacy risks which must be addressed.

McPherson, R., Jana, S., & Shmatikov, V. (2015). "No Escape from Reality: Security and Privacy of Augmented Reality Browsers." Presented at the the 24th International Conference on World Wide Web, 743-753.

Cyberbullying: What’s the Problem?

Canada’s current federal, provincial, and municipal definitions and understandings of cyberbullying are at odds with recent research. Cyberbullying is complex.In fact, the definition of cyberbullying depends on who is being asked. For example, public health, education, academic, and legal justice experts all differ in how they define the problem, measure prevalence, suggest causes, and identify the consequences of cyberbullying. Ultimately, inconsistent cyberbullying definitions make it difficult to accurately assess the problem, make decisions and measure policy performance. An uninformed and inconsistent definition can cause problems for the future of Canadian society.

Deschamps and McNutt from discuss and outline the current problems Canada faces in how it defines and combats cyberbullying at the federal, provincial, and municipal levels of government. Cyberbullying discussions have been increasing since 2005, but it was not until 2012 that cyberbullying debate really took off following the suicide of 3 Canadian teenagers (Amanda Todd, Rehtaeh Parsons, and Todd Loik). Public discussions of cyberbullying have most often focused on criminal incidents and legislation, but other government responses to the problem have involved various definitions of the problem. Differences in the definition of cyberbullying have led to a mixed bag of policy instruments and applications. Some government approaches have involved legal institutions, while others have used education and prevention campaigns to address cyberbullying incidents. Additionally, some have relied on administrative institutions where responses to online abuse involved the use of electronic filters or restricted access to computers. Ultimately, there is a need to form a common measurement of what classifies a certain behaviour as cyberbullying. This is also needed within academia and research, as numerous studies respond and define the problem differently. Without a consistent approach to the issue, policy makers will continue to struggle to provide effective measures for cyberbullying within Canada.

Contrary to popular belief, cyberbullying is not just traditional face-to-face bullying with an added element of computer communication. Definitions implying that cyberbullying is similar to traditional face-to-face bullying will not provide decision-makers with the right information needed to implement effective policies and legislation. Policy makers must recognize that cyberbullying has unique features that are distinct from traditional face-to-face bullying. Cyberbullying is not simply a new interpretation of an old problem, but a new phenomenon entirely.

There are 3 main flaws with the current understanding of cyberbullying. The first flaw is the assumption that cyberbullying starts in school. While there are instances of cyberbullying that do originate in school, schools are not the sole site of cyberbullying. Social media and cell phone mobility allows people to stay continuously connected to their online worlds. Consequently, cyberbullying goes beyond schools and enters personal life. The second flaw of the current definition of cyberbullying is the assumption that youths are the main victims of cyberbullying. While bullying among adults is less frequent than among youth, it is still not uncommon. Studies show that 10% to 15% of adult workers experience some form of bullying. The final flaw of the current definition is that it assumes that cyberbullying behaviour is always intentional. Recent research suggests that there are developmental differences in how youths morally evaluate bullying activities. Harmful cyberbullying can be unintentional. Ultimately, having the dominant understanding of cyberbullying built on traditional approaches to bullying restricts policy makers from properly treating and responding to instances of cyberbullying.

Governmental decision-making needs a research-based, well-informed, standardized definition of cyberbullying as a basis for effective policies and legislation.

Deschamps, R., & McNutt, K. (2016). “Cyberbullying: What’s the problem?” Canadian Public Administration, 59 (1), 45-71.

Fast, Lean, and Accurate: Modeling Password Guessability Using Neural Networks

The most common way of identifying yourself online is using a password. People are not very good at choosing passwords. This makes guessing passwords easier than it could be. Systems that evaluate our password choices can help us to select better passwords. Unfortunately, the software for thorough password checking can be very complex, requiring time and resources that are impractical. This means that the password checkers we use currently are forced to be overly simplistic and inaccurate.

Melicher et al. propose using artificial neural networks to guess passwords. Artificial neural networks are a machine-learning technique that models human neurons. They are well suited to generating text and approximate classifications. It makes sense that they would be well suited to guessing passwords too, but this had never been tested.

The researchers experimented with different methods of training and implementing neural networks to guess passwords and to rate the 'guessability' or strength of passwords. They developed a neural network method that appears to be better than other methods, particularly when targeting more difficult passwords and when making more guesses. Despite this, they note it is still a good idea to use multiple methods together for better estimates of password strength. The neural networks method used only a tiny fraction of the disk space required by methods. The method can also be compressed further so that can be downloaded as part of a webpage. This is advantageous as it allows prospective passwords to be gauged without them ever travelling across a network. Used this way, the meter was able to measure password resistance to guessing more precisely than models currently in use.

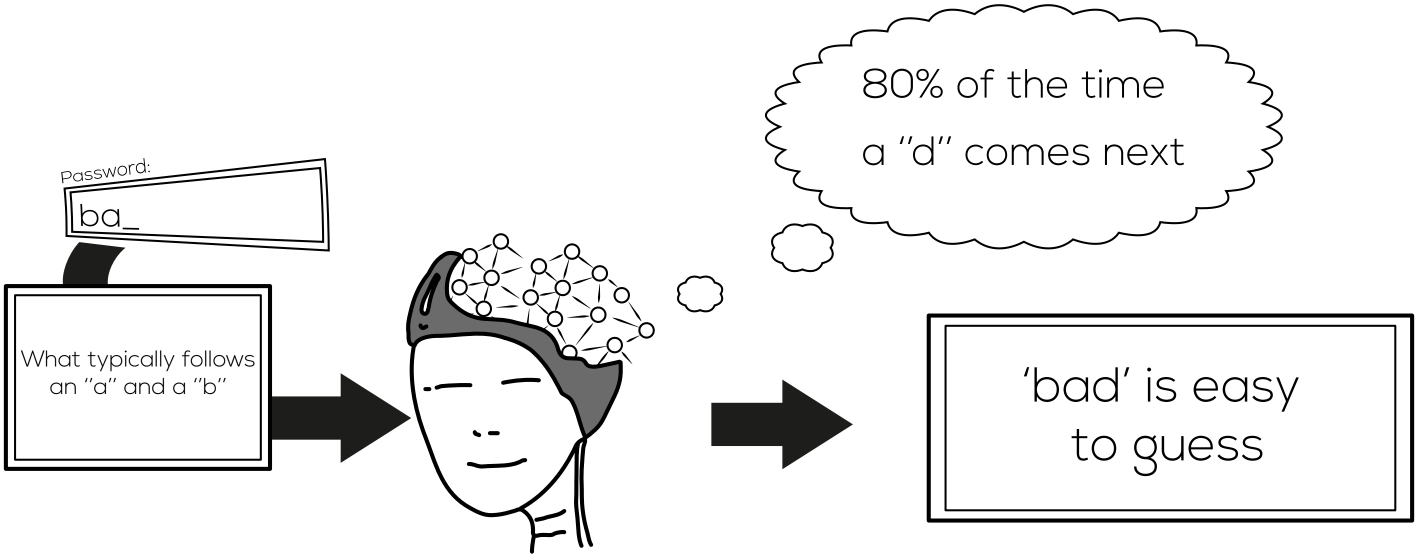
Password:
ba_
What typically follows an "a" and a "b"
80% of the time a "d" comes next
'bad' is easy to guess

Using neural networks to model password choices and measure their strength is not only possible, but it also offers benefits over current approaches. Password strength meters are only part of the puzzle of improving password selection and do not offer a total solution. Nonetheless, neural network based password selection assistance could prove to be valuable to improving security.

Artificial Neural Networks could provide valuable password strength checking tools.

Melicher et al. (2016). "Fast, Lean, and Accurate: Modeling Password Guessability Using Neural Networks." Presented at the 25th USENIX Security Symposium.

SERENE-RISC Six Key Activities

Smart Cybersecurity Network (SERENE-RISC) is a knowledge mobilization network that organises six key activities intended to reach its various audiences: workshops and seminars, a knowledge brokers' forum, quarterly knowledge digests, Konnect - online knowledge-sharing platform, a public website and a professional development program.

Workshops & Seminars

October 2016
Ottawa

Knowledge Brokers

Expanded Access Program

Knowledge Digest

Sponsorship
Opportunities Available

Konnect

Nearly 700 hand-selected resources on cybersecurity including exclusive content

Website

Cybersecurity tips section, news on the network and Digest archive.

Professional Development

Ask us about the Graduate Development Sessions

The SERENE-RISC Quarterly Cybersecurity Knowledge Digest

2016 Summer
Editor-in-Chief Michael Joyce; **Scientific Editor** Sonia Chiasson
Editor: Shannon McPhail
Schilder, J. D. et al. summarised by Raushan Ara Dilruba sponsored by SERENE-RISC
Chachra, N. et al summarised by Xi Liu sponsored by SERENE-RISC
McPherson, R. et al. summarised by Yomna Abdelaziz sponsored by SERENE-RISC
Deschamps R., & McNutt, K. summarised by Jin Lee sponsored by SERENE-RISC

Links to external sources for papers are provided for convenience only and do not represent an endorsement or guarantee of the external service provider.

Government of Canada
Networks of Centres of Excellence
Gouvernement du Canada
Réseaux de centres d'excellence

serene risc

Université de Montréal