Min-Entropy as a Resource

Geoffrey Smith
Florida International University

ISR, 15 October 2011, Oxford, UK

Secrecy is crucial to security goals

- **In itself**: Anonymity protocols want to keep originator identities secret.

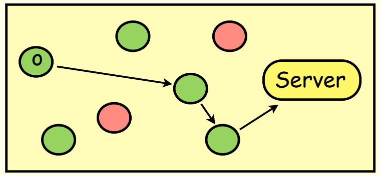

- **As a tool**: Operating systems need to keep passwords secret to achieve authentication.

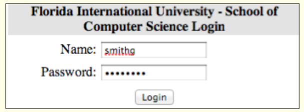

But leaks are hard to avoid

Password checking

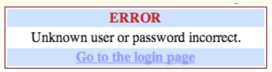

Statistical queries

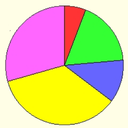

Timings of decryptions

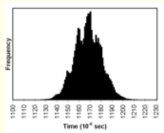

Quantitative Information Flow

- An active area of research for the past decade [ClarkHuntMalacaria02, ...]
- Suppose a system gets a secret input S and produces an observable output O.
- Want to model "how much" information about S is leaked to an adversary $\mathcal{A}$ who sees O.
- Can we view the secrecy of S as a "resource" that is gradually "consumed" by the system?
- A first, clearcut, example: O = S & 0x007f;
  - If S is a 32-bit integer, and all 2^{32} values are equally likely, then this program leaks 7 bits (out of 32) to O.

A subtler example: Crowds Protocol [RubinReiter98]

- Users wish to communicate anonymously with a server.
- The originator first sends the message to a randomly-chosen forwarder (possibly itself).
- Each forwarder forwards it again with probability p_f, or sends it to the server with probability $1-p_f$.
- But some crowd members are collaborators that report who sends them a message.
- Some information about the originator may be leaked. But how much???

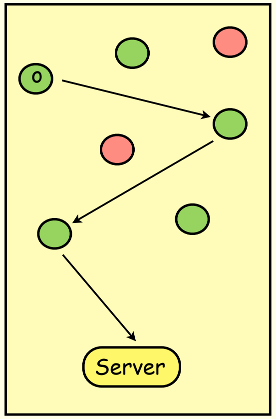

Plan of the talk

- Motivation
- Min-entropy leakage
  - Bayes vulnerability, min-entropy, min-capacity
  - Basic properties
- Consumption of min-entropy in composed channels
  - Channel composition operators
  - Bounds on min-entropy leakage of composed channels
  - Application to timing attacks on cryptography
- The dynamic perspective on leakage

Measuring secrecy

- Assume S is a random variable with distribution P_S.
- Assume, for the worst case, that P_S is known to the adversary $\mathcal{A}$.
- Initially, how "secret" or "uncertain" is S to $\mathcal{A}$?
- Shannon entropy [1948] is a classic measure:
  - $H(S) = -\Sigma_s P_S[s] \log P_S[s]$
- But this does not work so well for secrecy.
  - If $P_S = (1/2, 2^{-1000}, 2^{-1000}, 2^{-1000}, ..., 2^{-1000})$, then $H(S) = 500.5$ bits.
  - But half the time $\mathcal{A}$ can guess S correctly in one try!

Bayes Vulnerability and Min-Entropy

- [Smith09] proposed to focus instead on S's Bayes vulnerability to be guessed by $\mathcal{A}$ in one try, and to measure secrecy using min-entropy [Rényi61]:
- **Definition**: $V(S) = \max_s P_S[s]$
- **Definition**: $H_\infty(S) = -\log V(S)$
- If $P_S = (1/2, 2^{-1000}, 2^{-1000}, 2^{-1000}, 2^{-1000}, ..., 2^{-1000})$, then $V(S) = 1/2$ and $H_\infty(S) = 1$ bit.
- [Indeed, the same is true if $P_S = (1/2, 1/2)$.]

Systems as information-theoretic channels

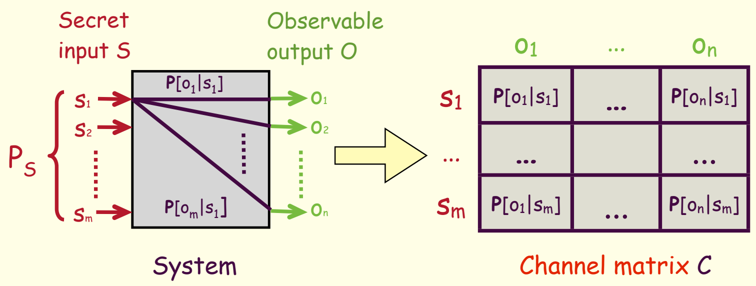

Each row of channel matrix C sums to 1.
C is deterministic if each entry is 0 or 1.

S has a priori distribution P_S .

Joint and a posteriori distributions

- Multiplying row s of C by $P_S[s]$ gives the joint matrix $P[s,o] = P_S[s]C[s,o]$
- By marginalization, we get a random variable O with distribution $P[o] = \ \Sigma_s P[s,o]$.
- **Bayes' Theorem**: $P[s|o] = P[o|s]P[s]/P[o] = P[s,o]/P[o]$
- So for each value o of O, we also get an a posteriori distribution $P_{S|o}$ by normalizing column o of the joint matrix.
- Assuming that $\mathcal{A}$ knows C and P_S, the distribution $P_{S|o}$ is what $\mathcal{A}$ knows about S if it sees output o.

An example channel and its distributions

Channel matrix

0	0	1/3	2/3
0	4/5	0	1/5
4/7	0	2/7	1/7
4/9	0	4/9	1/9

Joint matrix

0	0	1/16	1/8
0	1/4	0	1/16
1/8	0	1/16	1/32
1/8	0	1/8	1/32

$P_S = (3/16, 5/16, 7/32, 9/32)$ ← A priori distribution on S

$P_O = (1/4, 1/4, 1/4, 1/4)$ ← Distribution on O

A posteriori distributions on S:

$P_{S|o_1} = (0, 0, 1/2, 1/2)$
$P_{S|o_2} = (0, 1, 0, 0)$
$P_{S|o_3} = (1/4, 0, 1/4, 1/2)$
$P_{S|o_4} = (1/2, 1/4, 1/8, 1/8)$

Quantifying leakage

- S's initial secrecy is $H_\infty(S)$.
- Need to define S's remaining secrecy after $\mathcal{A}$ sees O.
- Intuitive equation:
  "leakage = initial secrecy – remaining secrecy"
- Clearly the "remaining secrecy" is based on the a posteriori distributions on S.
- But how should it be defined?

V(S|O) and $H_\infty(S|O)$

- We consider the average vulnerability, over all runs.
- **Definition**: $V(S|O) = \sum_o P[o]\, V(S|o)$
- $V(S|O) = \sum_o P[o] \max_s P[s|o] = \sum_o \max_s P[s,o]$
- V(S|O) is the complement of the Bayes risk. [ChatzikokolakisPalamidessiPanangaden08]
- Define $H_\infty(S|O)$, the remaining secrecy, as before.
- **Definition**: $H_\infty(S|O) = -\log V(S|O)$
  - Not defined by Rényi.
  - Not the same as $H_\infty(S|O) = \sum_o P[o]\, H_\infty(S|o)$.

V(S) and V(S|O) on example channel

Channel matrix

0	0	1/3	2/3
0	4/5	0	1/5
4/7	0	2/7	1/7
4/9	0	4/9	1/9

A priori distribution P_S (3/16, 5/16, 7/32, 9/32)

Joint matrix

0	0	1/16	1/8
0	1/4	0	1/16
1/8	0	1/16	1/32
1/8	0	1/8	1/32

- $V(S) = \max_s P_S[s] = 5/16$
- $V(S|O) = \sum_o \max_s P[s,o] = 1/8 + 1/4 + 1/8 + 1/8 = 5/8$
- S's Bayes vulnerability doubles.
- A priori, $\mathcal{A}$ guesses that S is s_2.
- A posteriori, $\mathcal{A}$'s best guess for S depends on O:

  $o_1 \to s_3$ (or s_4), $o_2 \to s_2$, $o_3 \to s_4$, $o_4 \to s_1$

Min-entropy leakage

- **Definition**: Min-entropy leakage

  $$\mathcal{L}_{SO} = H_\infty(S) - H_\infty(S|O) = \log \frac{V(S|O)}{V(S)}$$

- So leaking x bits means increasing the Bayes vulnerability by a factor of 2^x.
- In the example, $\mathcal{L}_{SO} = \log 2 = 1$ bit.
- **Definition**: Min-capacity

  $\mathcal{ML}(C)$ is the maximum min-entropy leakage, over all a priori distributions P_S.

Properties of min-entropy leakage

- **Theorem**: $V(S|O) \geq V(S)$, so $\mathcal{L}_{SO} \geq 0$.
- **Theorem**: $\mathcal{ML}(C)$ is the log of the sum of the column maximums of C.
  - Also, $\mathcal{ML}(C)$ is realized by a uniform a priori P_S.
- **Corollary**: If C is deterministic, then $\mathcal{ML}(C)$ is the log of the number of feasible outputs.
- **Corollary**: $\mathcal{ML}(C) = 0$ iff the rows of C are identical.
- $\mathcal{L}_{SO} = 0$ if S and O are independent. Not conversely!
- Indeed $\mathcal{L}_{SO} = 0$ if O never affects $\mathcal{A}$'s best guess.

Example ("base-rate fallacy")

- Consider a good, but imperfect, test for cancer:

channel matrix

	positive	negative
cancer	0.90	0.10
no cancer	0.07	0.93

- A priori (age 40-50, no symptoms, no family history)

  P_S[cancer] = 0.008 P_S[no cancer] = 0.992

joint matrix

	positive	negative
cancer	0.00720	0.00080
no cancer	0.06944	0.92256

column maximums

- $V(S|O) = 0.992 = V(S)$, so $\mathcal{L}_{SO} = 0$.
- Always guess "no cancer"! (P[cancer|positive] ≈ 0.094)

Plan of the talk

- Motivation
- Min-entropy leakage
  - Bayes vulnerability, min-entropy, min-capacity
  - Basic properties
- Consumption of min-entropy in composed channels
  - Channel composition operators
  - Bounds on min-entropy leakage of composed channels
  - Application to timing attacks on cryptography
- The dynamic perspective on leakage

Channels in Cascade C_1C_2 [EspinozaSmith11]

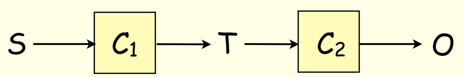

- Formed by multiplying two channel matrices, or by factoring a channel matrix.
- **Theorem**: In channel C_1C_2, for any P_S, $\mathcal{L}_{SO} \le \mathcal{L}_{ST}$.
  - Analogue of the data-processing inequality.
  - Curiously, we **can** have $\mathcal{L}_{SO} > \mathcal{L}_{TO}$.
- **Theorem**: $\mathcal{ML}(C_1C_2) \le \min \{ \mathcal{ML}(C_1), \mathcal{ML}(C_2) \}$
- **Corollary** [KöpfSmith10]: $\mathcal{ML}(C_1C_2) \le \log|\mathcal{T}|$
  - Here $\mathcal{T}$ is the set of feasible values for T.

Repeated independent runs $C^{(n)}$

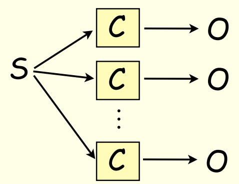

(Useful only when C is probabilistic!)

- $C^{(n)}[s, (o_1,o_2,\ldots,o_n)] = \prod_i C[s, o_i]$
- Curiously, even if $\mathcal{L}_{SO} = 0$, we **can** have $\mathcal{L}_{SO^2} > 0$.
- **Theorem** [BorealePampoloniPaolini11]:

  $\mathcal{ML}(C^{(n)})$ converges exponentially quickly to the log of the number of **distinct** rows in C.
  - Intuitively, distinct rows of the channel matrix can be distinguished by repeatedly sampling O.

More on $C^{(n)}$

- However, $\mathcal{ML}(C^{(n)})$ grows only logarithmically in n.
- **Theorem** [KöpfSmith10]: $\mathcal{ML}(C^{(n)}) \leq |\mathcal{O}| \log(n+1)$.
  - Here $\mathcal{O}$ is the set of feasible values of O.
  - The proof factors $C^{(n)}$ into the cascade of two channels with a small set $\mathcal{T}$ of intermediate values.
  - In fact we have $\mathcal{ML}(C^{(n)}) \leq \log\binom{n+|\mathcal{O}|-1}{n}$.
- [Could we get a stronger bound based on $\mathcal{ML}(C)$?]

Application: timing attacks on cryptography

- Remote timing attack [BonehBrumley03].
- 1024-bit RSA key recovered in 2 hours from standard OpenSSL implementation across LAN.

SSL Handshake (simplified)

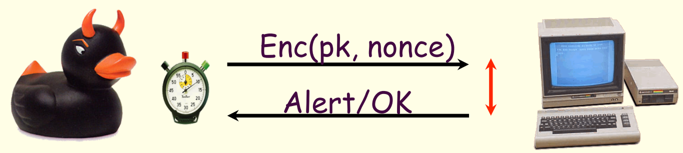

$\mathcal{A}$ can estimate the type to decrypt each nonce with secret key sk.

Effectiveness of blinding and bucketing against timing attacks [KöpfSmith10]

- Blinding: randomize ciphertext before decryption, and de-randomize after decryption.
- Bucketing: force decryption to take one of a small number of possible times.
  - Using as few as 5 buckets costs little performance.
- Thanks to blinding, we have a "repeated independent runs" channel, so the previous theorem applies.
- **Corollary**: With blinding, the min-capacity of the timing attack is logarithmic in the number of timing observations.
  - With 5 buckets (so $|\mathcal{O}|=5$) and 2^{40} timing observations, the min-capacity is at most 155.4 bits.

A more powerful composition $C_1 + C_2$

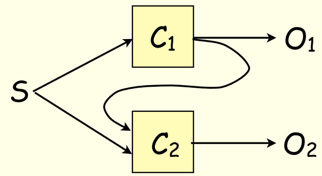

(Both O_1 and O_2 are output.)

- The "back arrow" allows adaptive processing.
  - C_2 can decide what to do based on the output of C_1.
  - Write $C_1 +_{na} C_2$ for composition without the back arrow.
- **Theorem** [BartheKöpf11]:

  $\mathcal{ML}(C_1 + C_2) \leq \mathcal{ML}(C_1) + \mathcal{ML}(C_2)$
- [Contrast with $\mathcal{ML}(C_1C_2) \leq \min\{ \mathcal{ML}(C_1), \mathcal{ML}(C_2) \}$.]

Examples: Some deterministic channels with two feasible outputs (giving min-capacity of 1 bit).

- $E_i \equiv$ if (S == c_i) O = 1; else O = 0;

  $\mathcal{ML}(E_1 + E_1 + ... + E_n) \leq \log(n+1)$ bits
- $G_i \equiv$ if (S >= c_i) O = 1; else O = 0;

  $\mathcal{ML}(G_1 +_{na} G_2 +_{na} ... +_{na} G_n) \leq \log(n+1)$ bits

  But can have $\mathcal{ML}(G_1 + G_2 + ... + G_n) = n$ bits

```
if (S >= 512) O1 = 512; else O1 = 0;
if (S >= O1 + 256) O2 = 256; else O2 = 0;
if (S >= O2 + 128) O3 = 128; else O3 = 0;
...
```

Invariant:
$O_i \leq S < O_i + 2^{10-i}$

- $A_i \equiv$ if (S & 2^{i-1}) O = 1; else O = 0;

  $\mathcal{ML}(A_1 +_{na} A_2 +_{na} ... +_{na} A_n) = n$ bits.

Plan of the talk

- Motivation
- Min-entropy leakage
  - Bayes vulnerability, min-entropy, min-capacity
  - Basic properties
- Consumption of min-entropy in composed channels
  - Channel composition operators
  - Bounds on min-entropy leakage of composed channels
  - Application to timing attacks on cryptography
- The dynamic perspective on leakage

The dynamic perspective on leakage

- So far, we have considered the static perspective of leakage averaged over all runs.
- The dynamic perspective instead considers one particular run of C, producing a particular output o.
- In this case $\mathcal{A}$ can refine the distribution on S from P_S to $P_{S|o}$.
- In the earlier example, seeing o_2 shows $\mathcal{A}$ that S must be s_2, since $P_{S|o_2} = (0, 1, 0, 0)$.
- Moreover, if $\mathcal{A}$ can run the channel repeatedly, using the same value of S each time, then it can repeatedly refine its distribution on S.

Repeated refinement example

- With example channel, if we run repeatedly and observe outputs o_3, o_1, and o_3, then we have
- $P_S = (3/16, 5/16, 7/32, 9/32)$

  $P_{S|o_3} = (1/4, 0, 1/4, 1/2)$

  $P_{S|o_3o_1} = (0, 0, 9/23, 14/23)$

  $P_{S|o_3o_1o_3} = (0, 0, 81/277, 196/277)$
- $H_\infty(S) = -\log(5/16) \approx 1.678$

  $H_\infty(S|o_3) = -\log(1/2) = 1$

  $H_\infty(S|o_3o_1) = -\log(14/23) \approx 0.716$

  $H_\infty(S|o_3o_1o_3) = -\log(196/277) \approx 0.708$

A weakness of the dynamic perspective

- A particular run of a password checker could lead to total loss of the secret.
- How could we decide whether to allow *C* if it could lead to a total loss?
- Aborting execution in a bad case might itself reveal a lot of information!
- And we obviously need to distinguish between

  O = S;

  and

  if (S == c_i) O = 1; else O = 0;

Another weakness of the dynamic perspective

- More critically, min-entropy resulting from a particular run need not decrease monotonically!
- Suppose P_S = (9/10, 1/40, 1/40, 1/40, 1/40) and

C =

1	0
0	1
0	1
0	1
0	1

$V(S) = 9/10$	$H_\infty(S) \approx 0.152$
$V(S\|o_2) = 1/4$	$H_\infty(S\|o_2) = 2$

- [Example scenario: A medical test that refutes the only likely diagnosis.]
- So, under the dynamic perspective, min-entropy does not seem to behave as a reasonable "resource".

Conclusion

- Min-entropy can be viewed as a resource, and its leakage as a measure of the consumption of secrecy.
- Future directions:
  - Can min-entropy leakage be calculated for large systems?
  - How do min-entropy leakage and differential privacy fit together?
  - Min-entropy leakage is purely information theoretic. Could computational limits be incorporated?
- Thanks to my collaborators:

  Catuscia Palamidessi, Miguel Andrés, Boris Köpf, Ziyuan Meng, Barbara Espinoza

Some references

- Rényi, "On measures of entropy and information", 1961.
- Clark, Hunt, and Malacaria, "Quantitative analysis of the leakage of confidential data", ENTCS 2002.
- Köpf and Basin,"An information-theoretic model for adaptive side-channel attacks", CCS 2007.
- Chatzikoklakis, Palamidessi, Panangaden, "On the Bayes risk in information-hiding protocols", JCS 2008.
- Smith, "On the foundations of quantitative information flow", FOSSACS 2009.
- Braun, Chatzikokolakis, and Palamidessi, "Quantitative notions of leakage for one-try attacks", MFPS 2009.
- Köpf and Smith, "Vulnerability bounds and leakage resilience of blinded cryptography under timing attacks", CSF 2010.
- Boreale, Pampaloni, and Paolini, "Asymptotic information leakage under one-try attacks", FOSSACS 2011.
- Barth and Köpf, "Information-theoretic bounds for differentially private mechanisms", CSF 2011.
- Espinoza and Smith, "Min-entropy leakage of channels in cascade", FAST 2011.
- Smith, "Quantifying information flow using min-entropy", QEST 2011.

Discussion?