

# An Analysis of OpenSSL's Random Number Generator

## Eurocrypt 2016

Falko Strenzke



cryptosource GmbH,  
Darmstadt

[fstrenzke@cryptosource.de](mailto:fstrenzke@cryptosource.de)

©Falko Strenzke, cryptosource GmbH 2016

September 14, 2016


# Pseudo Random Number Generation

- Software-based RNG's use pseudo random number generators (PRNGs)
- but are not PRNGs




# Pseudo Random Number Generation

- Software-based RNG's use pseudo random number generators (PRNGs)
- but are not PRNGs




# Pseudo Random Number Generation

- Software-based RNG's use pseudo random number generators (PRNGs)
- but are not PRNGs



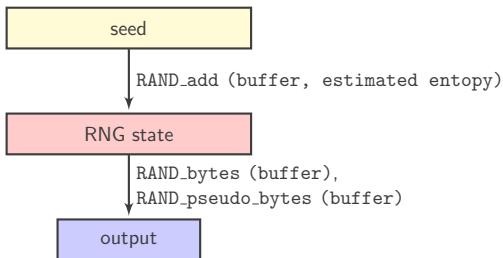
# Pseudo Random Number Generation

- Software-based RNG's use pseudo random number generators (PRNGs)
- but are not PRNGs

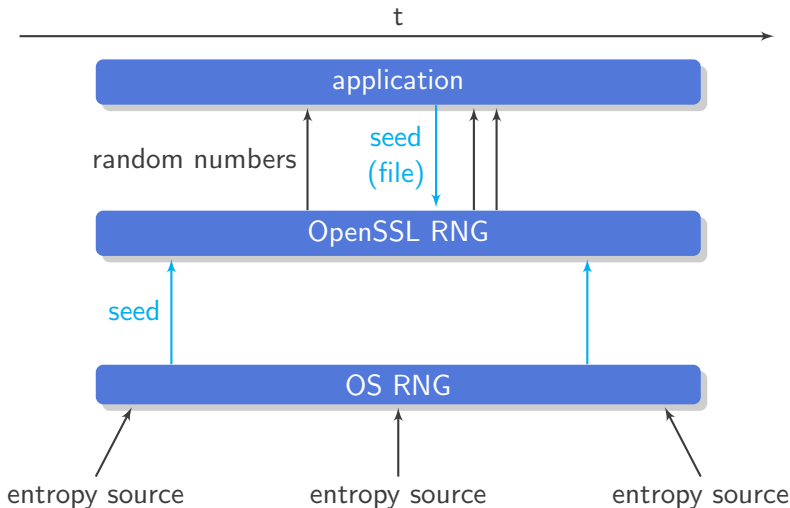


# Pseudo Random Number Generation

- Software-based RNG's use pseudo random number generators (PRNGs)
- but are not PRNGs




# Random Number Generation in Cryptographic Libraries



# Security Notions for RNGs

- forward security



- backward security




- don't leak any information about state in output




# Security Notions for RNGs

- forward security




- backward security




- don't leak any information about state in output

# Security Notions for RNGs

- forward security



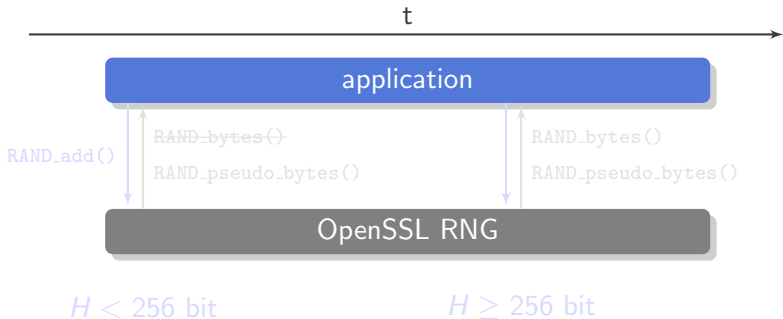
- backward security



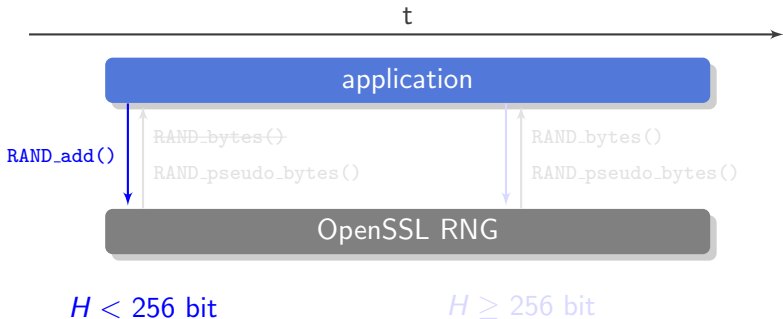
- don't leak any information about state in output

# Low Entropy Secret Leakage

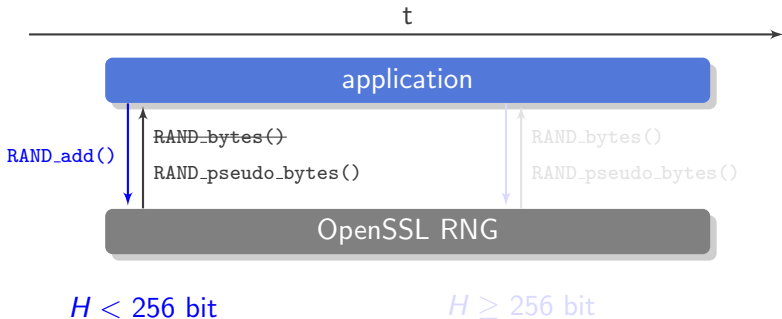
# Seeding by the Application



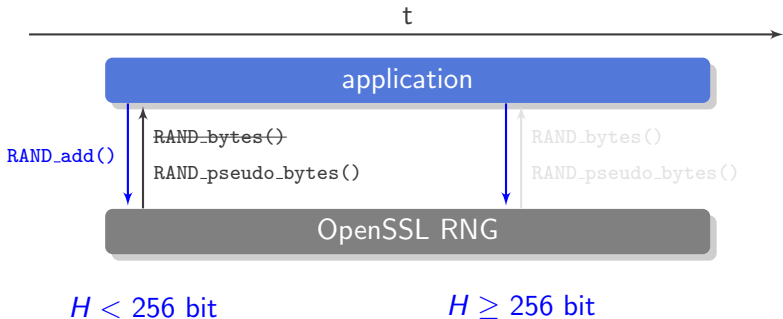
# Seeding by the Application



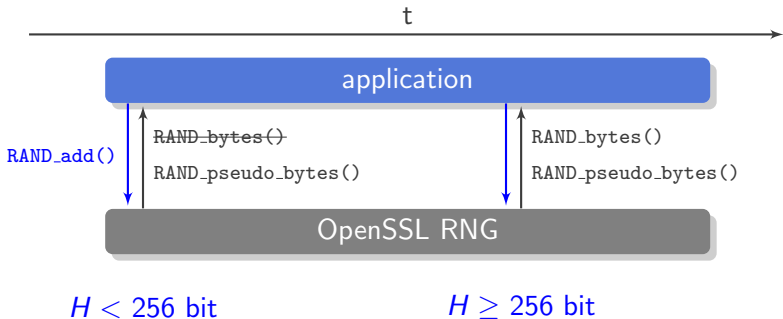
# Seeding by the Application



# Seeding by the Application

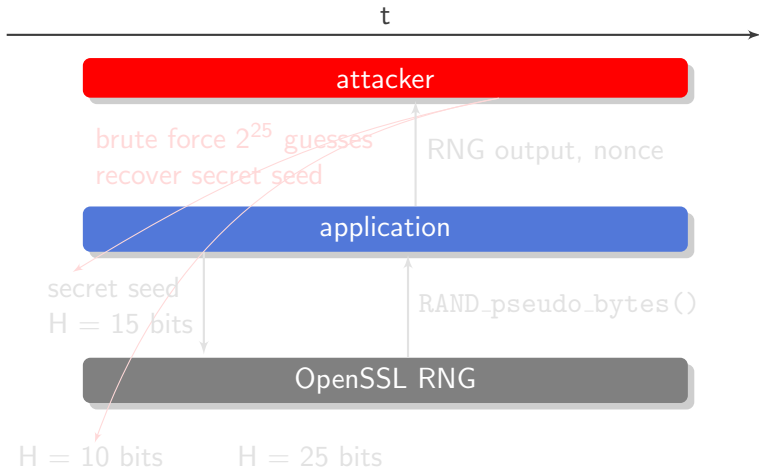


# Seeding by the Application

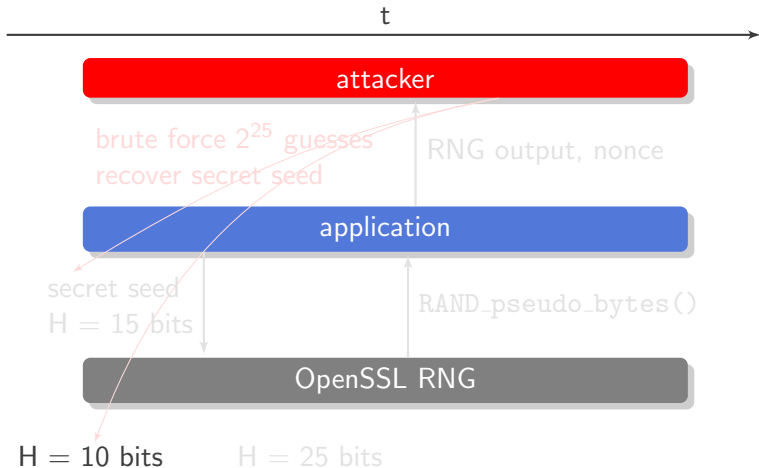




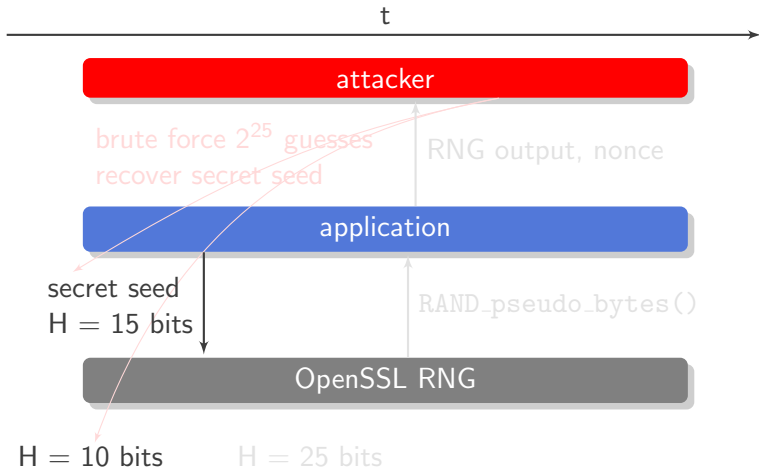
# Outputting Random Numbers in Low Entropy States



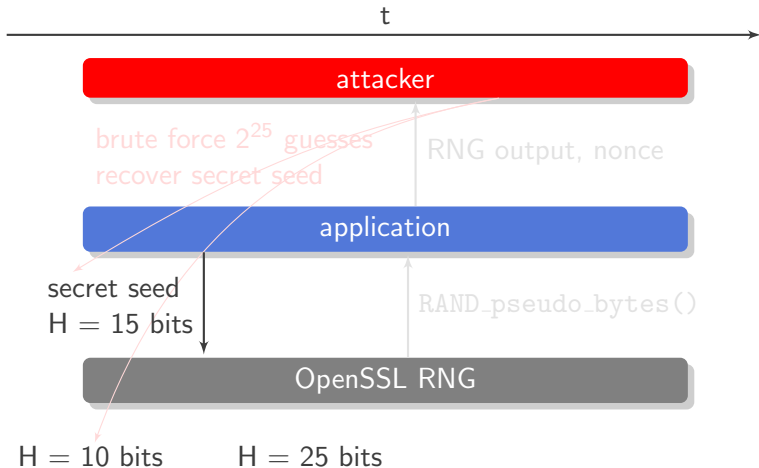
# Outputting Random Numbers in Low Entropy States



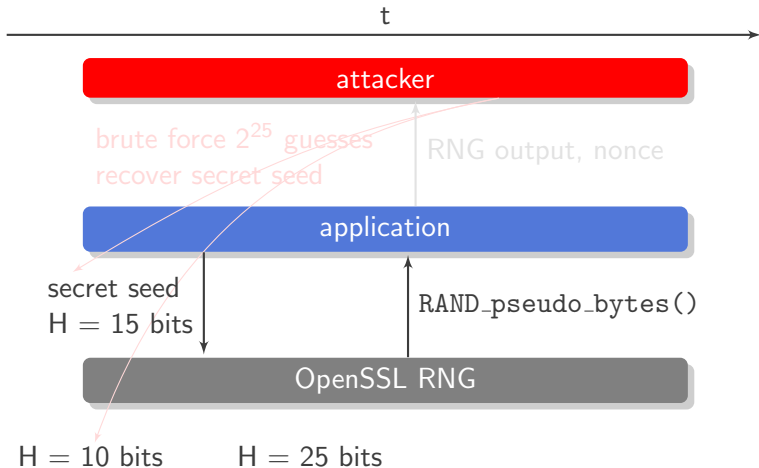
# Outputting Random Numbers in Low Entropy States



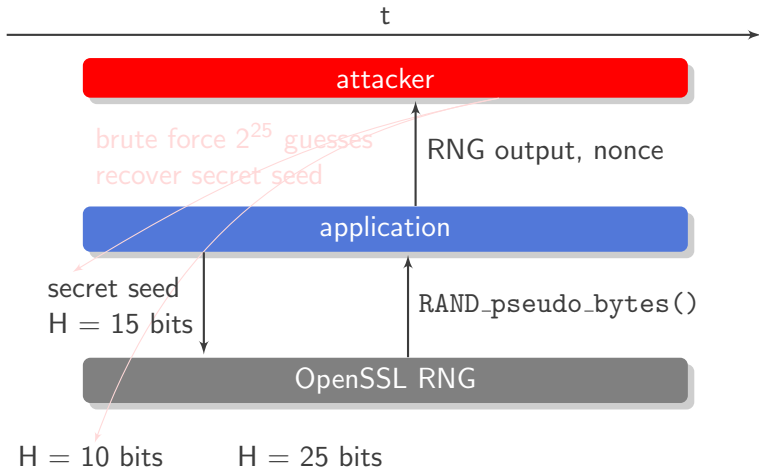
# Outputting Random Numbers in Low Entropy States



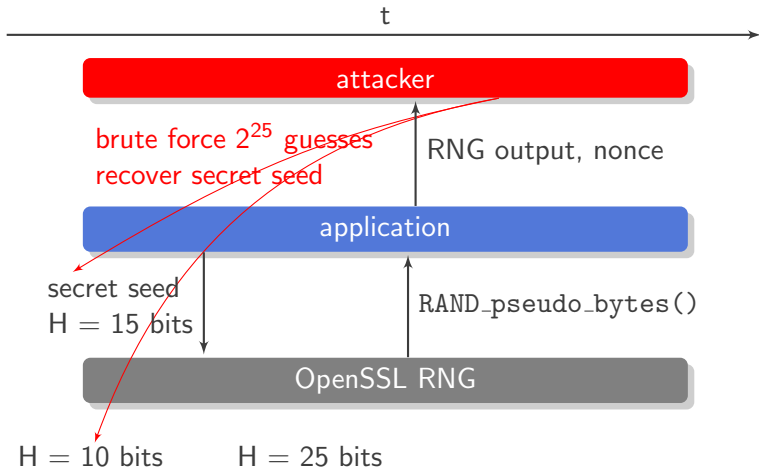
# Outputting Random Numbers in Low Entropy States



# Outputting Random Numbers in Low Entropy States



# Outputting Random Numbers in Low Entropy States



# Potentially Leaked Secrets

- `RAND_pseudo_bytes` generates output in the same way as `RAND_bytes`
- API documentation suggests to feed low-entropy secrets such passwords
- OpenSSL feeds the previous contents of buffers to be randomized to RNG state (Debian issue in 2008)
- previous contents could contain low entropy secrets by themselves
- overwriting secrets with random numbers is an established practice
- overwritten low entropy secrets may be leaked in output



# Potentially Leaked Secrets

- `RAND_pseudo_bytes` generates output in the same way as `RAND_bytes`
- API documentation suggests to feed low-entropy secrets such as passwords
- OpenSSL feeds the previous contents of buffers to be randomized to RNG state (Debian issue in 2008)
- previous contents could contain low entropy secrets by themselves
- overwriting secrets with random numbers is an established practice
- overwritten low entropy secrets may be leaked in output

# Potentially Leaked Secrets

- `RAND_pseudo_bytes` generates output in the same way as `RAND_bytes`
- API documentation suggests to feed low-entropy secrets such as passwords
- OpenSSL feeds the previous contents of buffers to be randomized to RNG state (Debian issue in 2008)
- previous contents could contain low entropy secrets by themselves
- overwriting secrets with random numbers is an established practice
- overwritten low entropy secrets may be leaked in output

# Potentially Leaked Secrets

- `RAND_pseudo_bytes` generates output in the same way as `RAND_bytes`
- API documentation suggests to feed low-entropy secrets such as passwords
- OpenSSL feeds the previous contents of buffers to be randomized to RNG state (Debian issue in 2008)
- previous contents could contain low entropy secrets by themselves
- overwriting secrets with random numbers is an established practice
- overwritten low entropy secrets may be leaked in output

- `RAND_pseudo_bytes` generates output in the same way as `RAND_bytes`
- API documentation suggests to feed low-entropy secrets such as passwords
- OpenSSL feeds the previous contents of buffers to be randomized to RNG state (Debian issue in 2008)
- previous contents could contain low entropy secrets by themselves
- overwriting secrets with random numbers is an established practice
- overwritten low entropy secrets may be leaked in output

- `RAND_pseudo_bytes` generates output in the same way as `RAND_bytes`
- API documentation suggests to feed low-entropy secrets such as passwords
- OpenSSL feeds the previous contents of buffers to be randomized to RNG state (Debian issue in 2008)
- previous contents could contain low entropy secrets by themselves
- overwriting secrets with random numbers is an established practice
- overwritten low entropy secrets may be leaked in output

# Core Cryptographic Function of OpenSSL's RNG

# Core Cryptographic Function of OpenSSL's RNG



- custom design
- ©1998

# Core Cryptographic Function of OpenSSL's RNG



- custom design
- ©1998




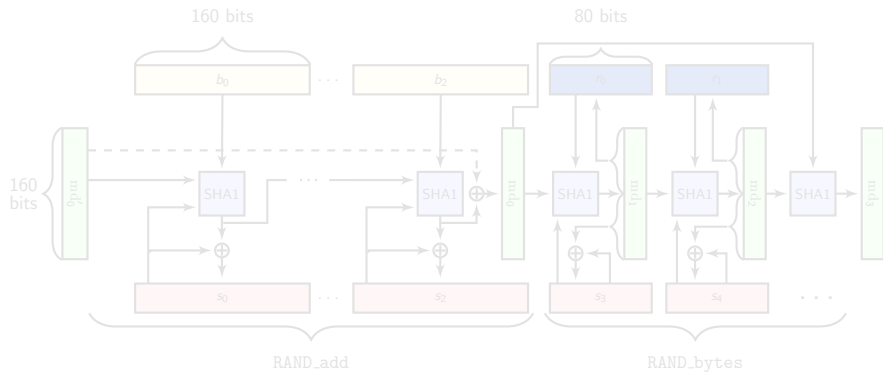
# Core Cryptographic Function of OpenSSL's RNG




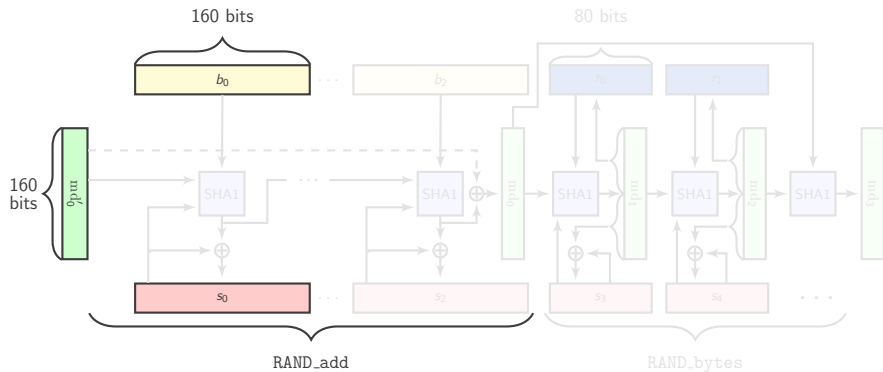
# Core Cryptographic Function of OpenSSL's RNG




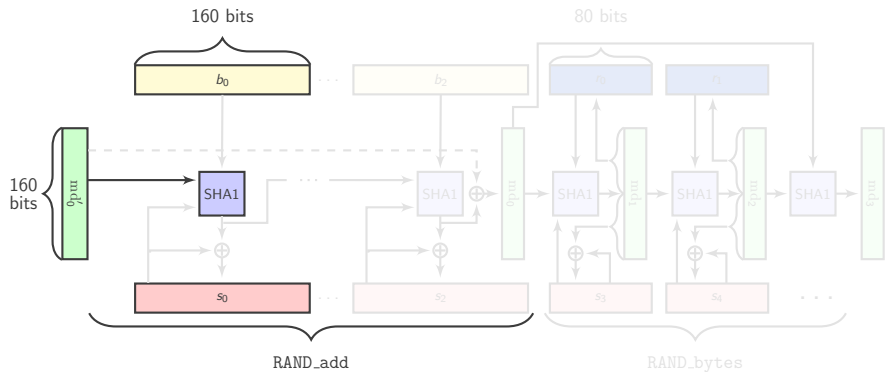
# Core Cryptographic Function of OpenSSL's RNG




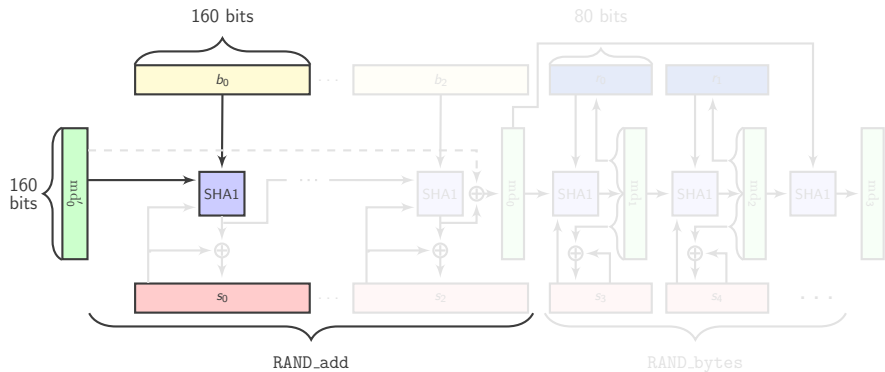
# Core Cryptographic Function of OpenSSL's RNG




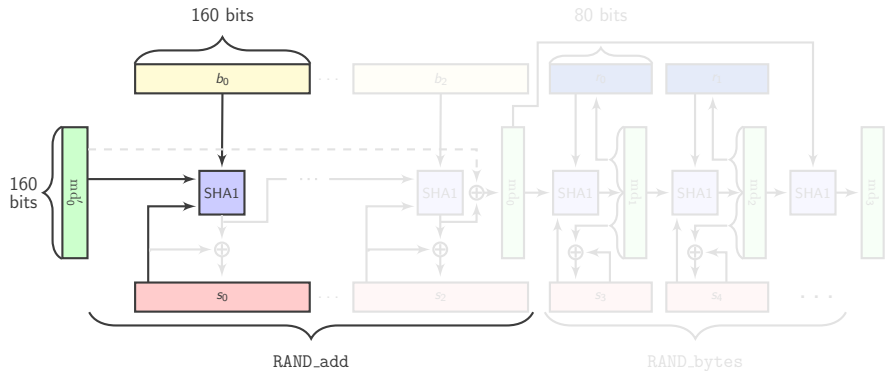
# Core Cryptographic Function of OpenSSL's RNG




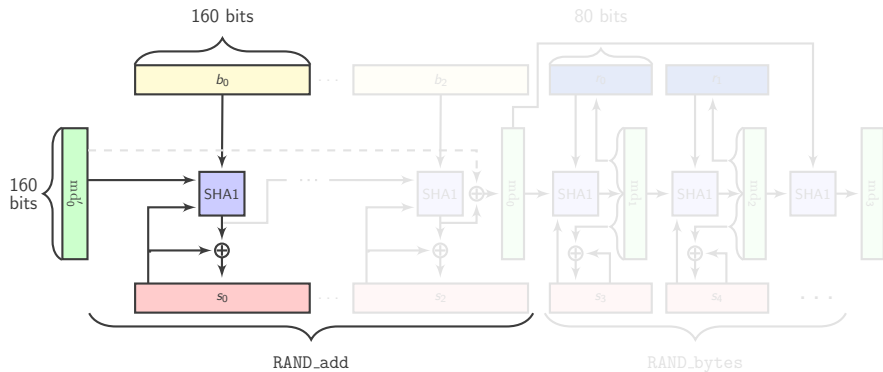
# Core Cryptographic Function of OpenSSL's RNG



# Core Cryptographic Function of OpenSSL's RNG


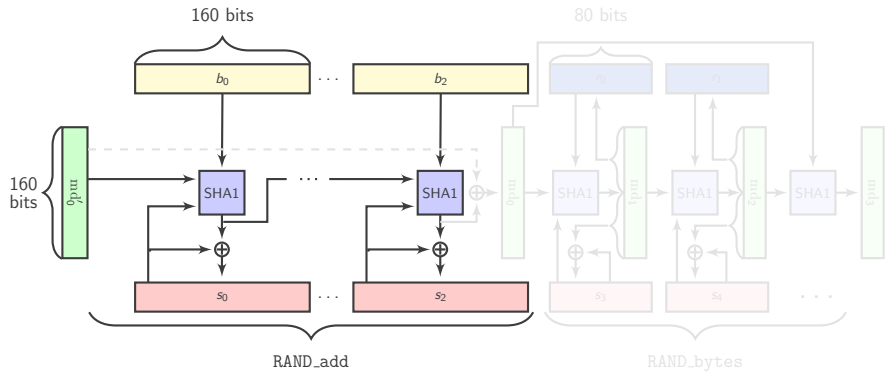


# Core Cryptographic Function of OpenSSL's RNG


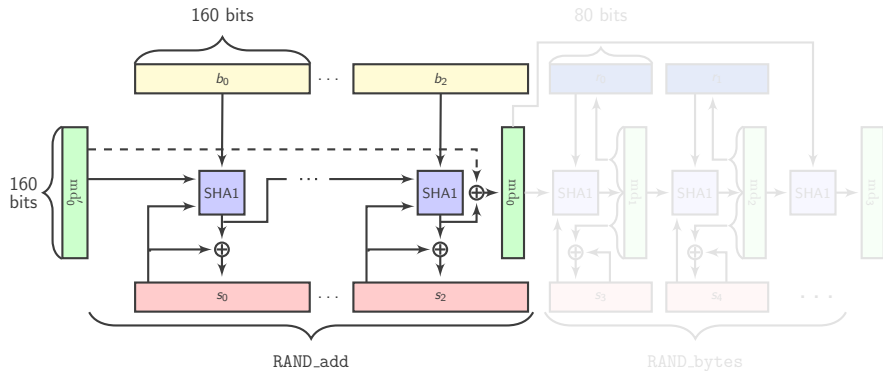





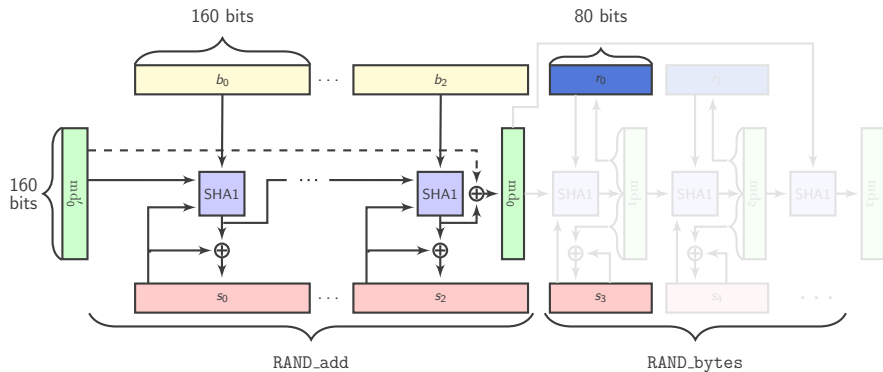
# Core Cryptographic Function of OpenSSL's RNG




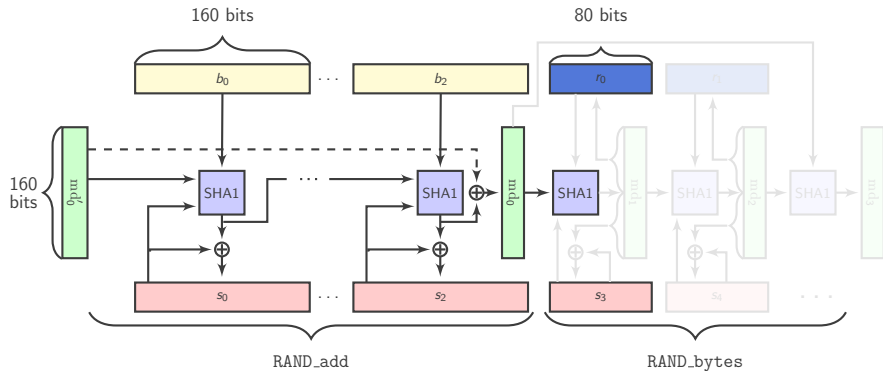
# Core Cryptographic Function of OpenSSL's RNG




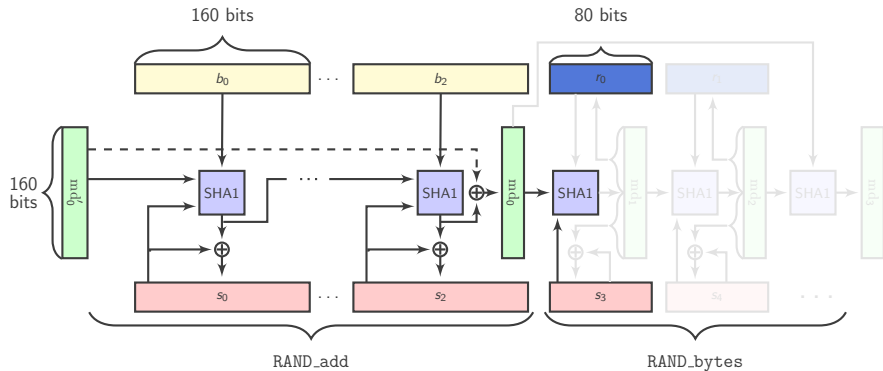
# Core Cryptographic Function of OpenSSL's RNG




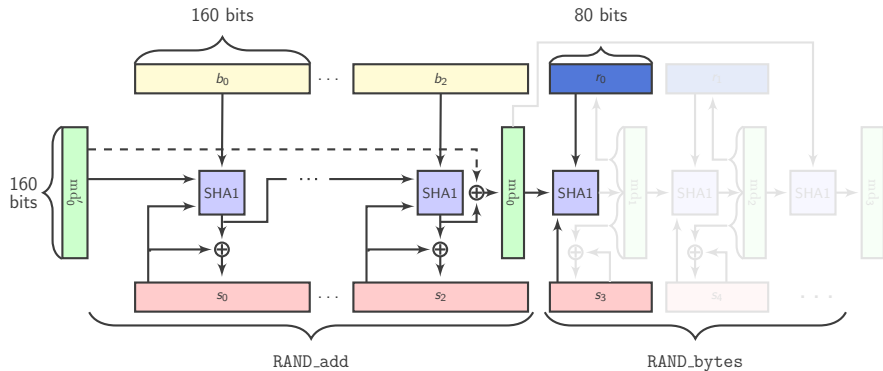
# Core Cryptographic Function of OpenSSL's RNG




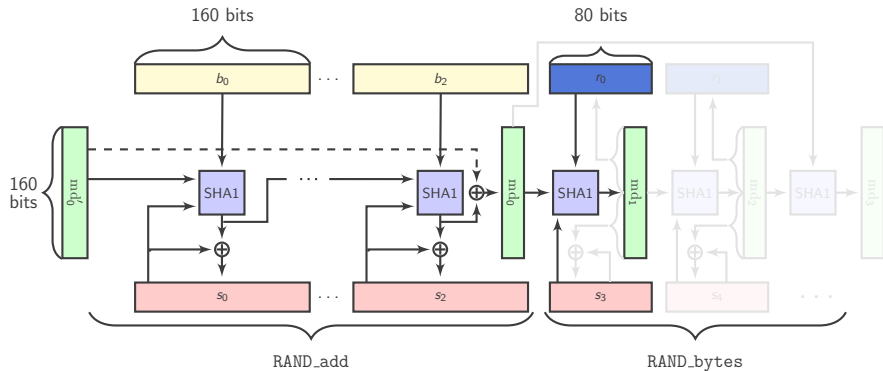
# Core Cryptographic Function of OpenSSL's RNG




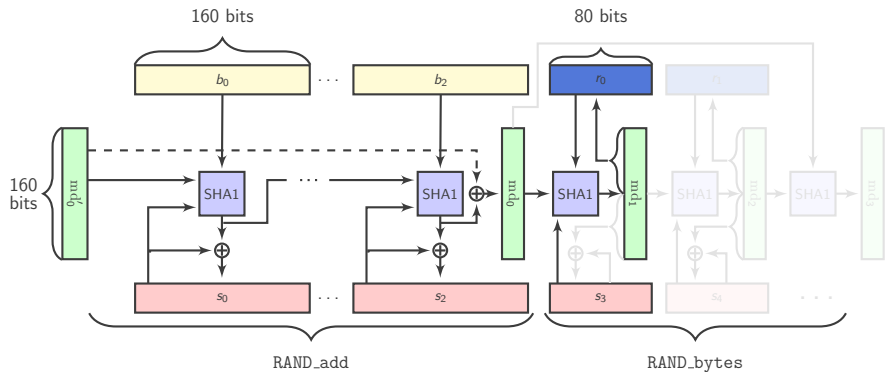
# Core Cryptographic Function of OpenSSL's RNG



# Core Cryptographic Function of OpenSSL's RNG


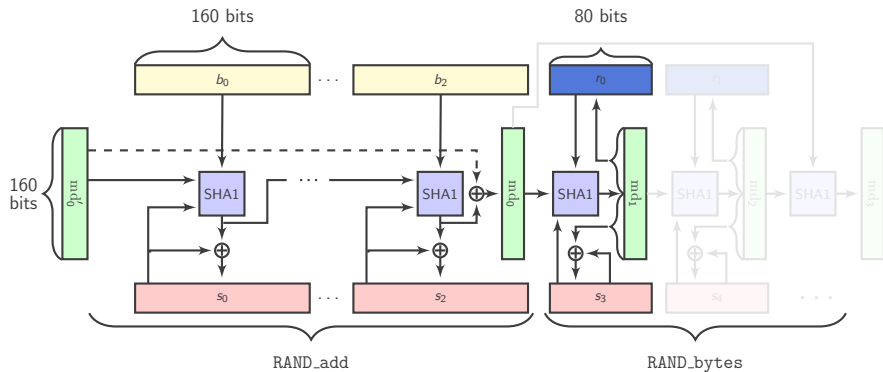


# Core Cryptographic Function of OpenSSL's RNG


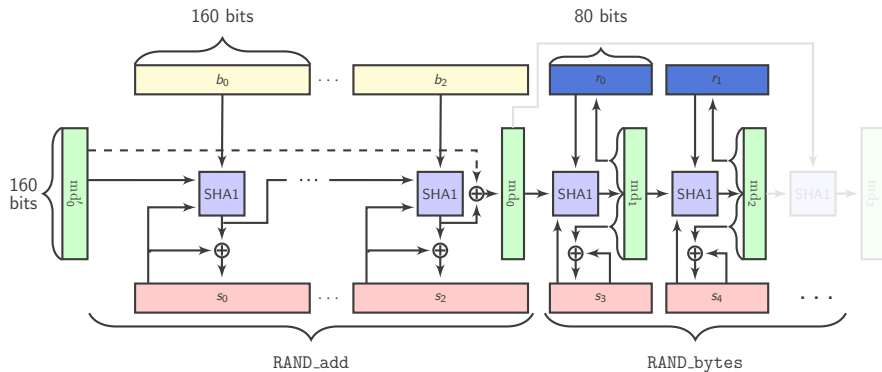





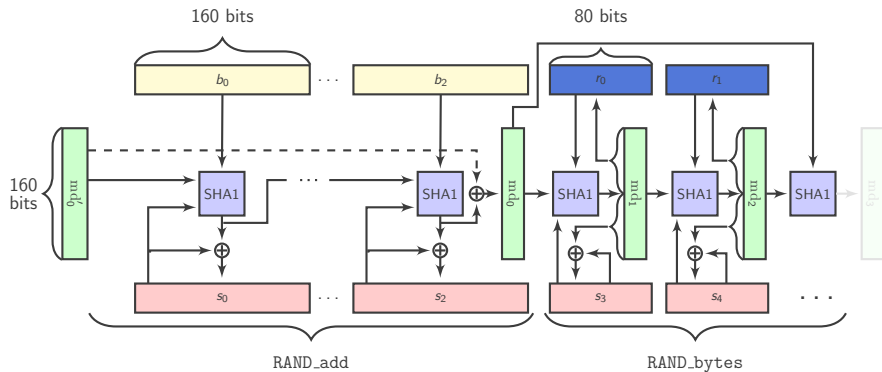
# Core Cryptographic Function of OpenSSL's RNG




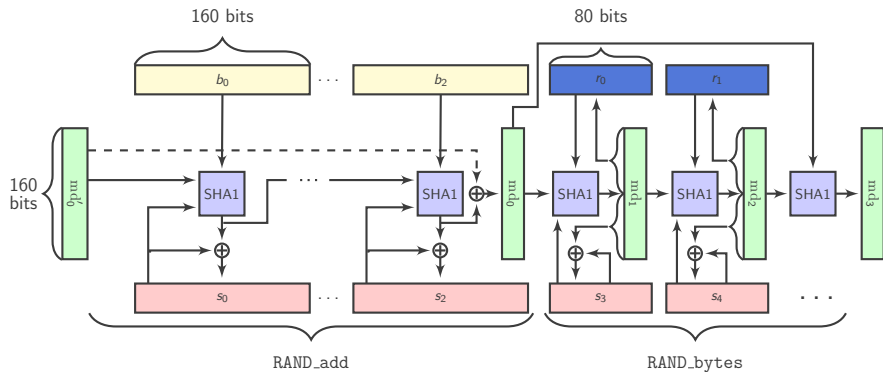
# Core Cryptographic Function of OpenSSL's RNG



# Core Cryptographic Function of OpenSSL's RNG

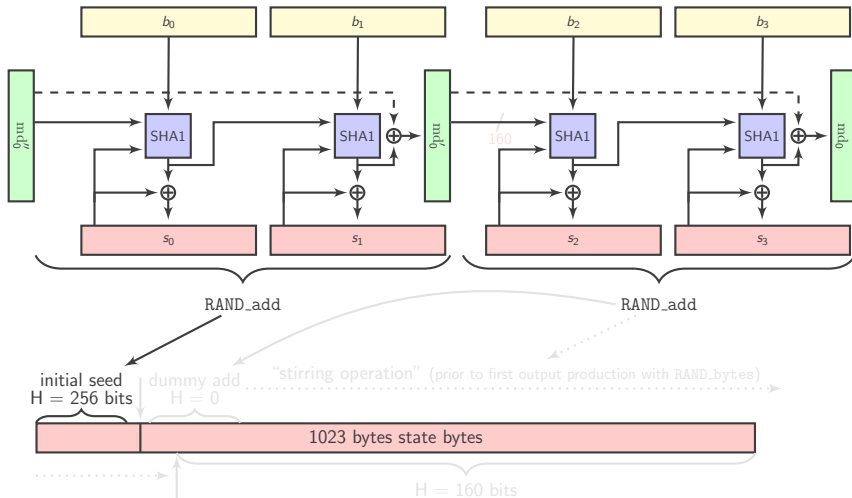


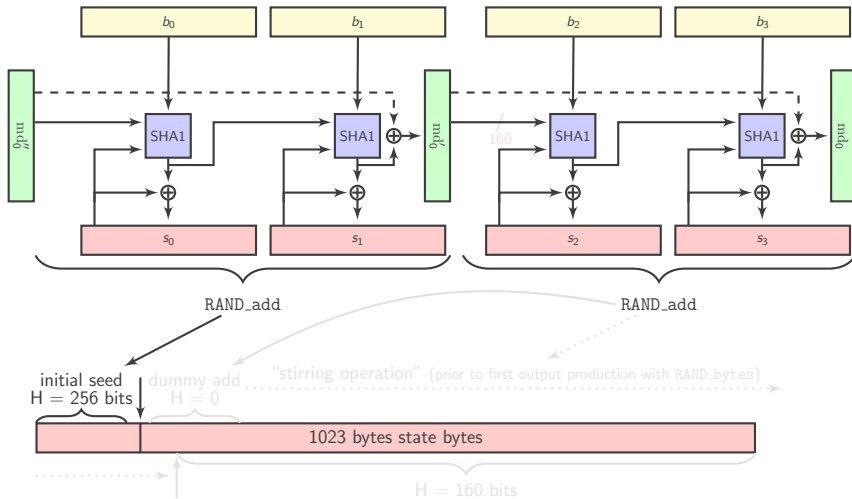
# Core Cryptographic Function of OpenSSL's RNG



# Output Entropy Limitation Vulnerabilities

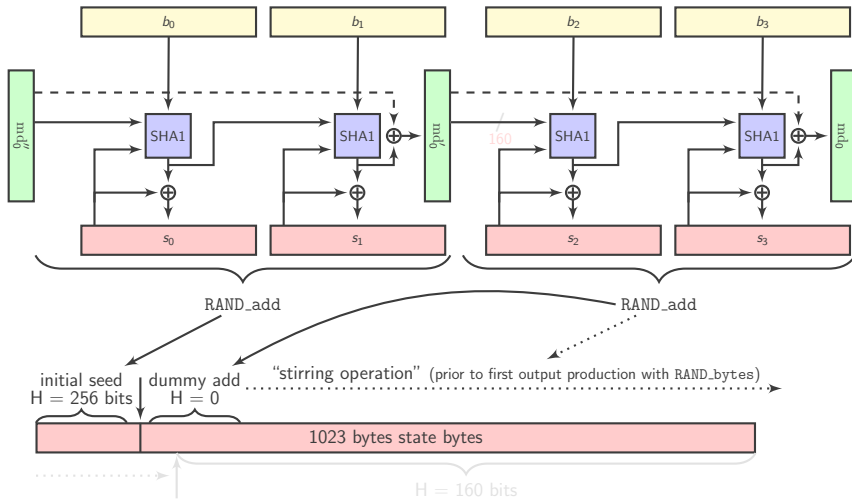


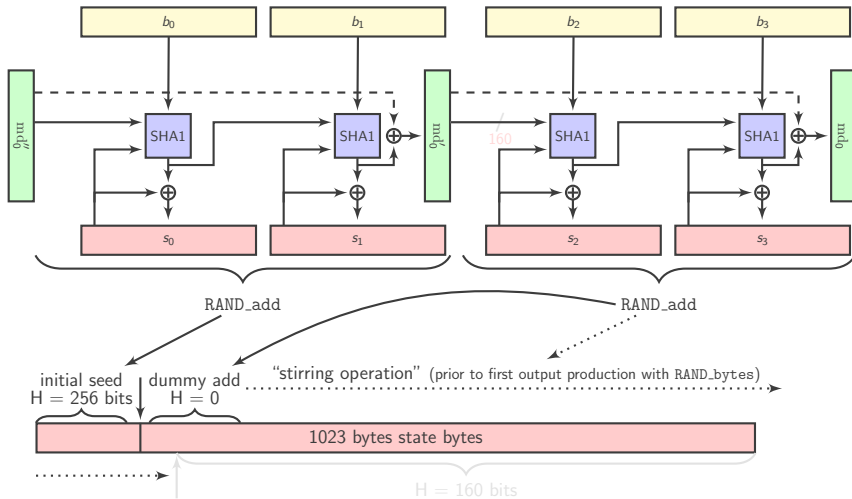




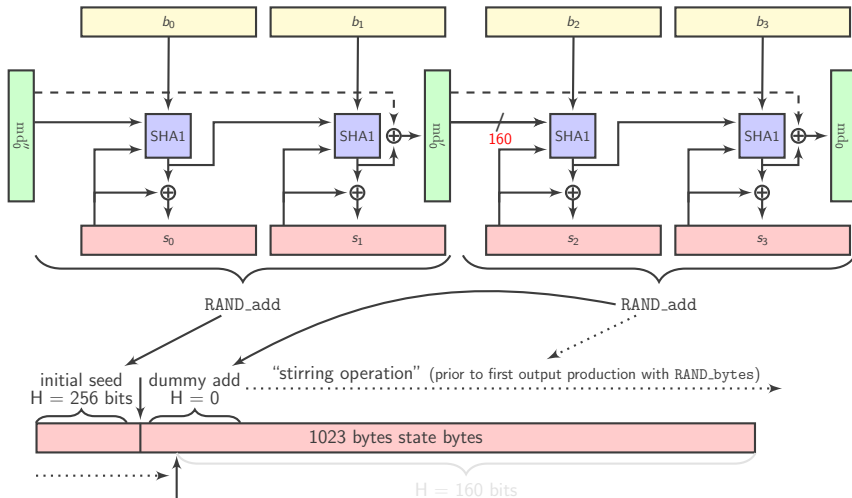


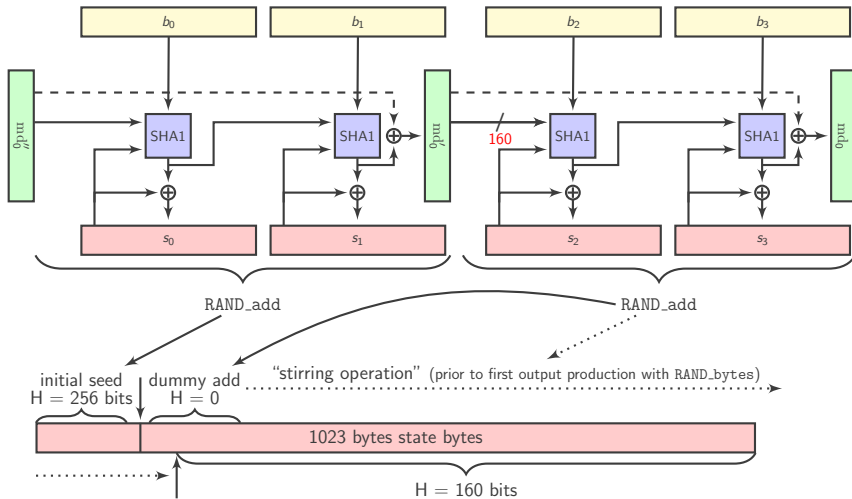


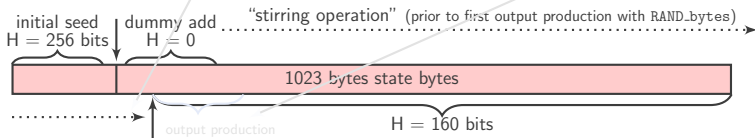
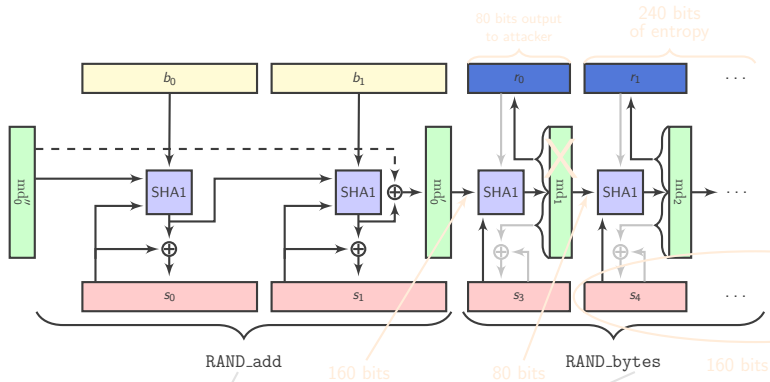


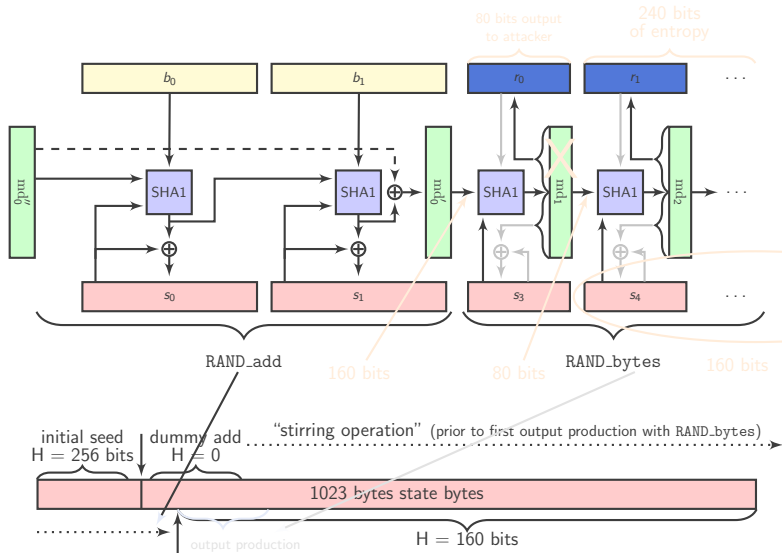




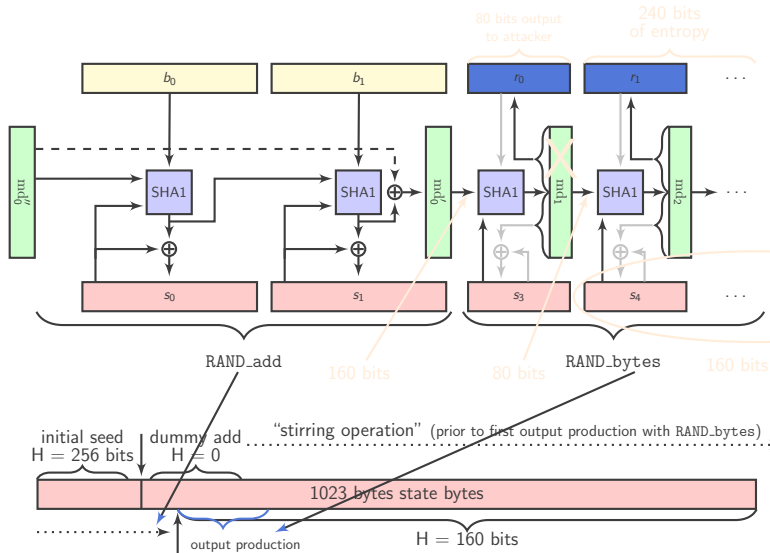


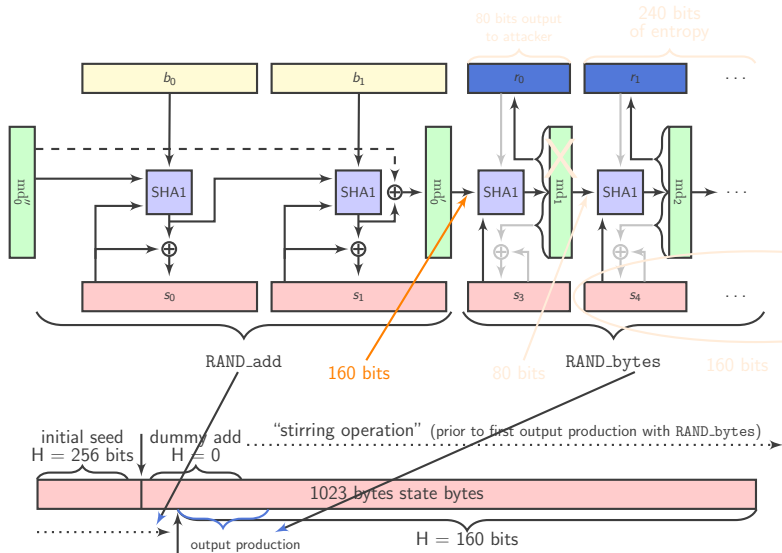


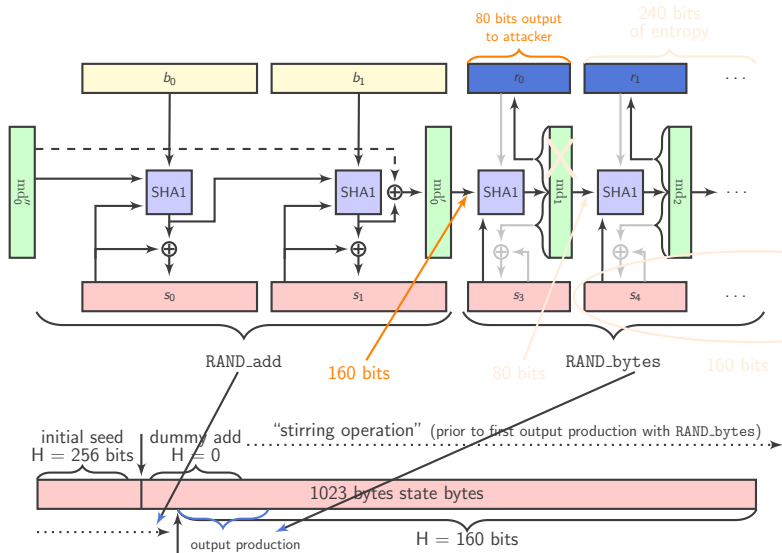


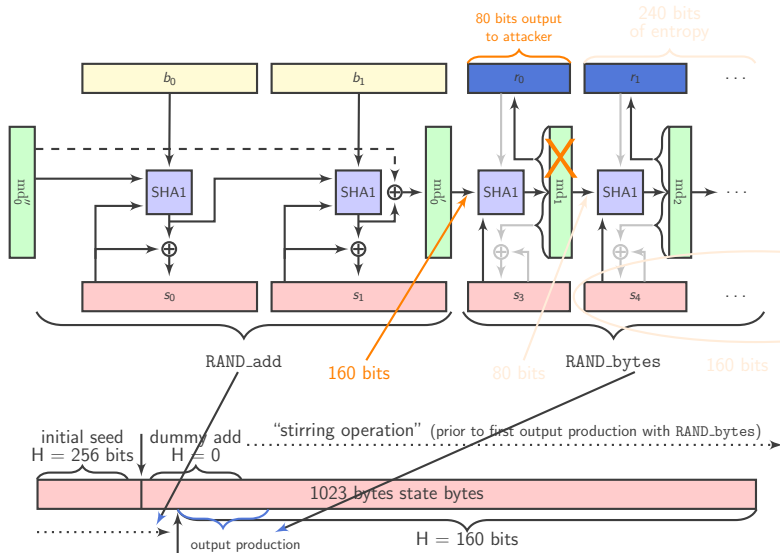


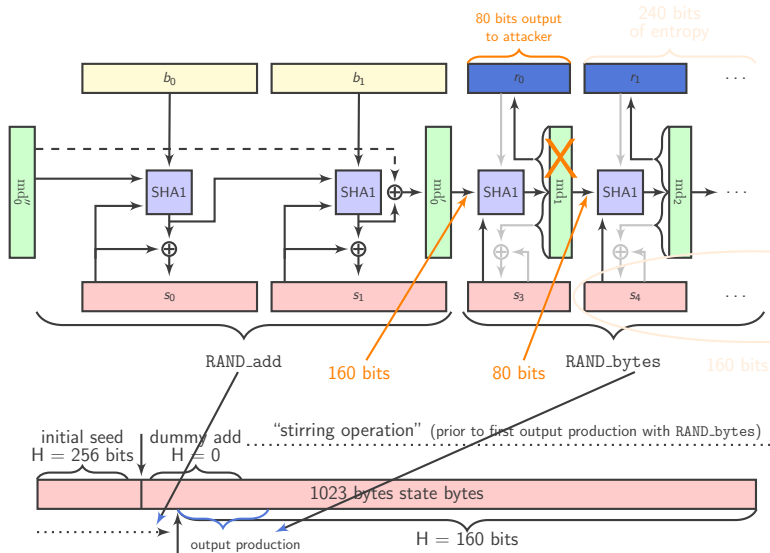



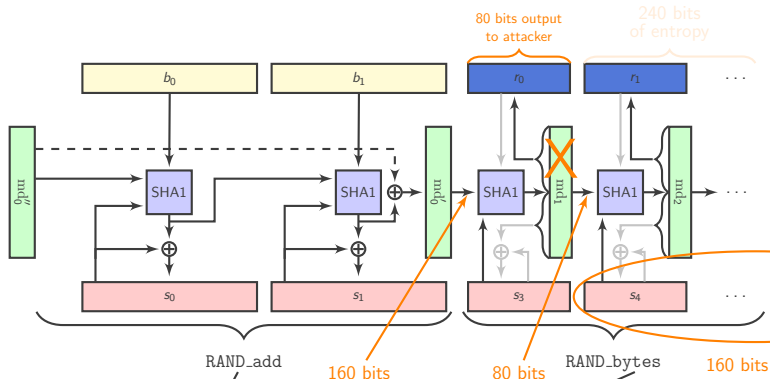



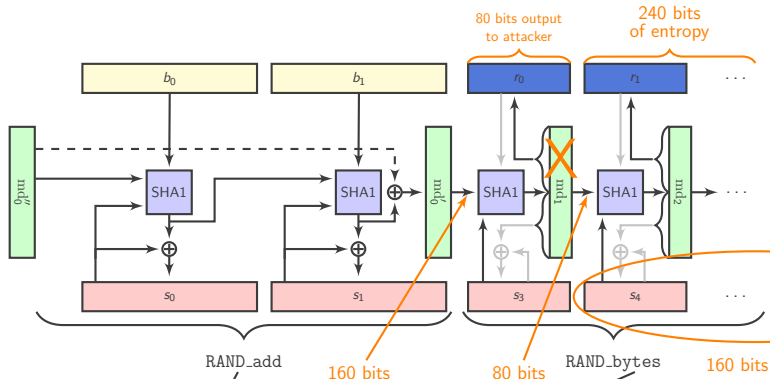




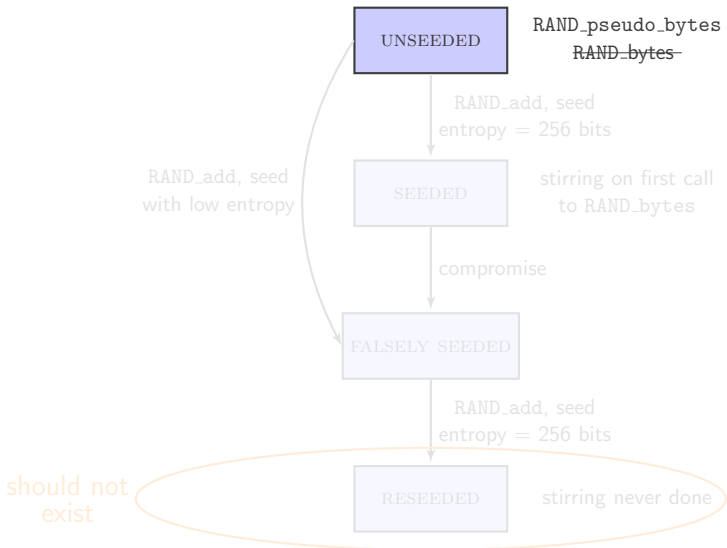






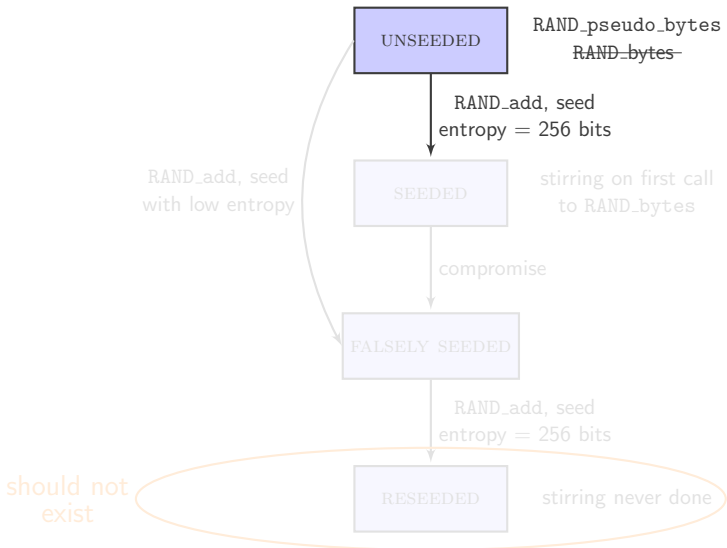


# Life cycle

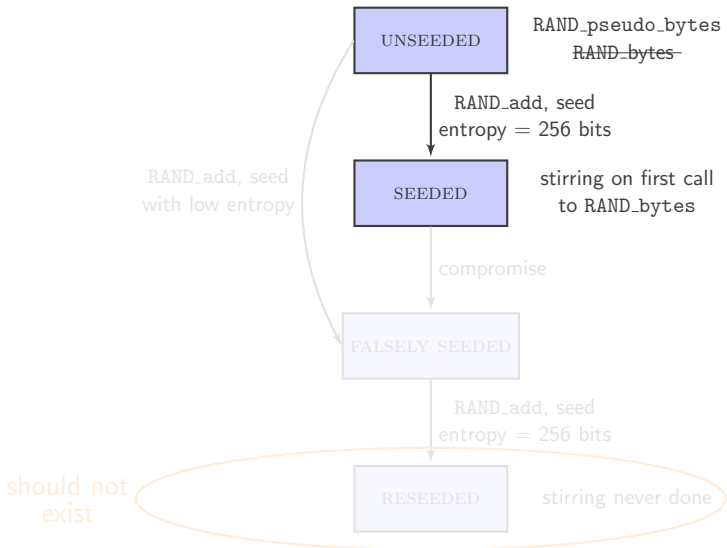




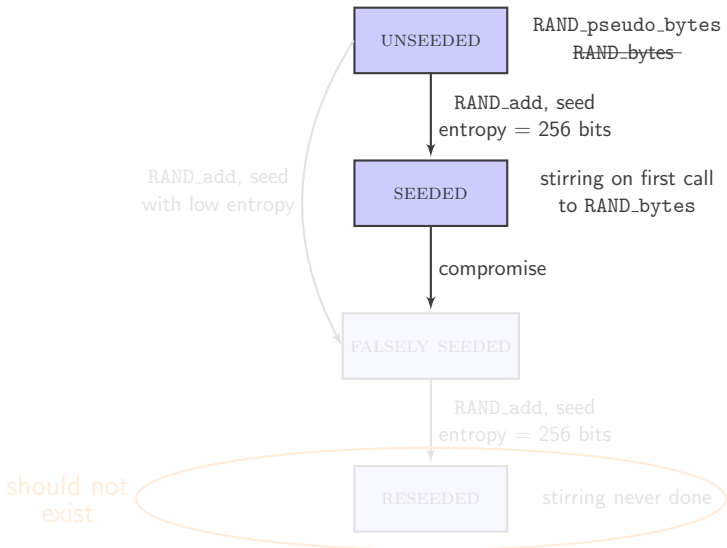
# Life cycle



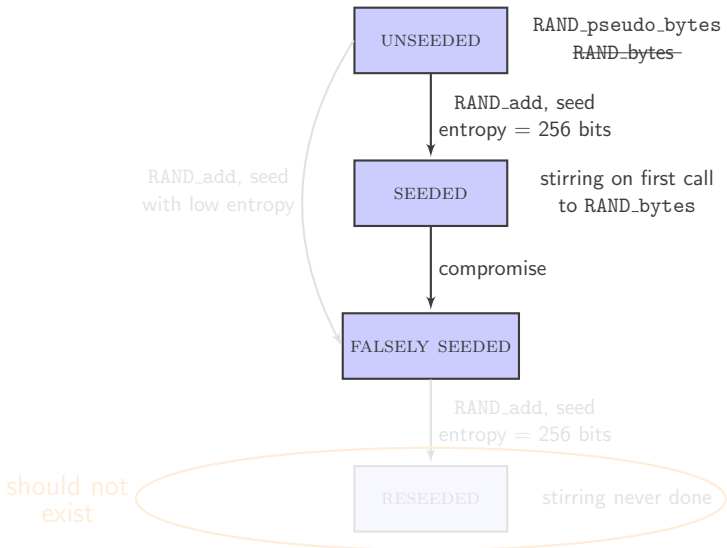
# Life cycle



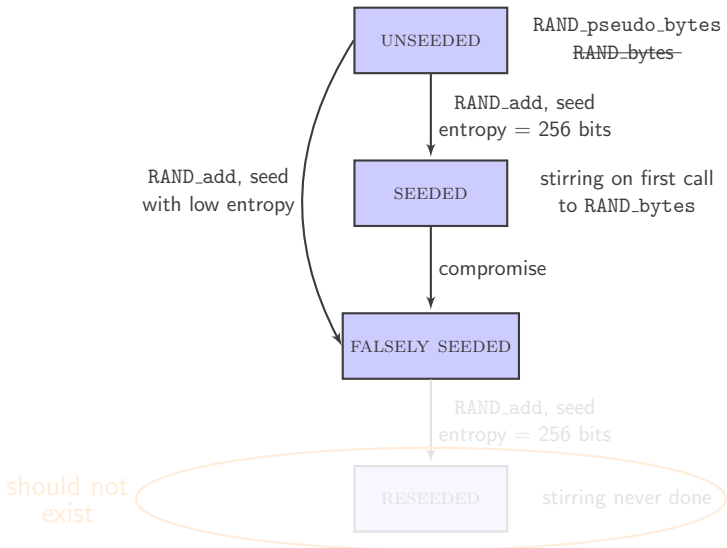
# Life cycle



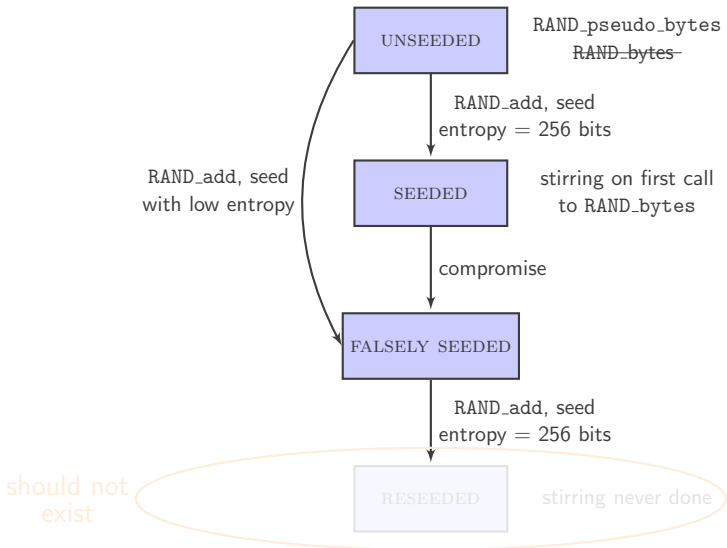
# Life cycle



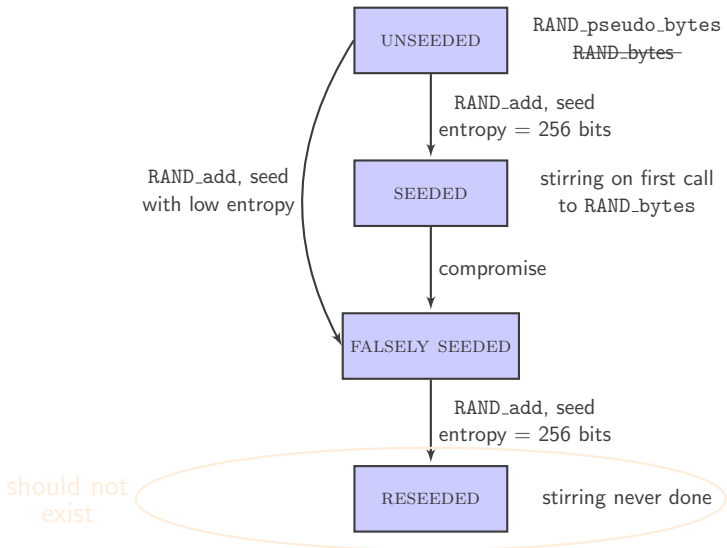
# Life cycle



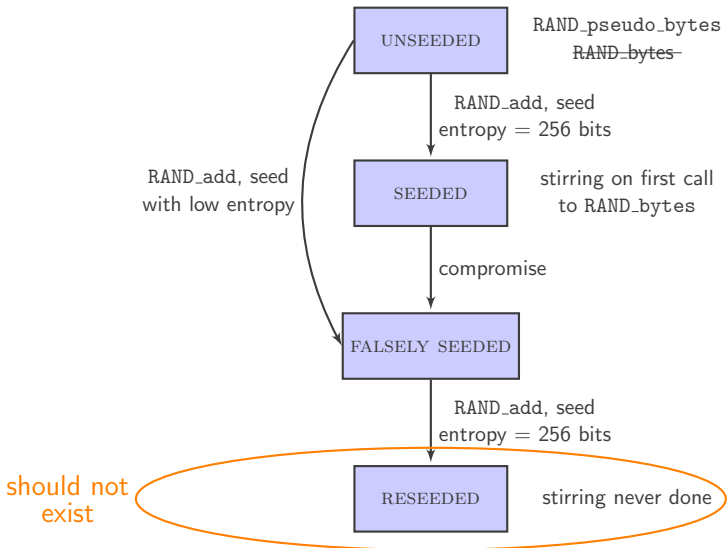
# Life cycle



# Life cycle

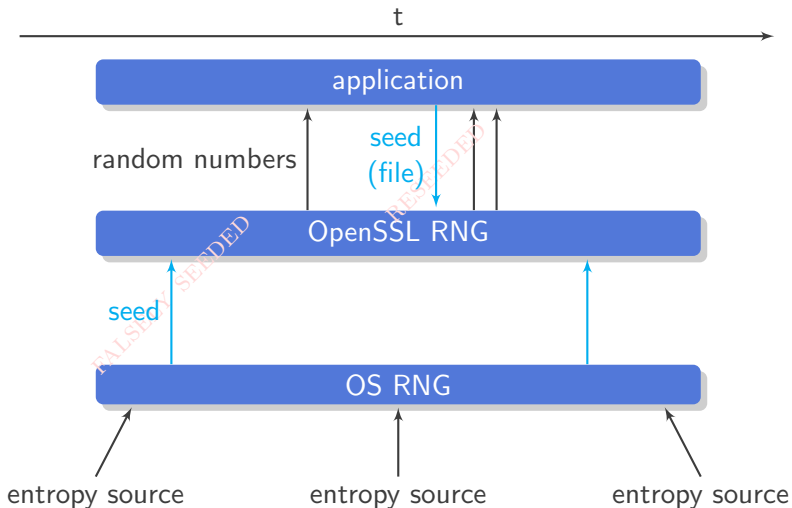


# Life cycle

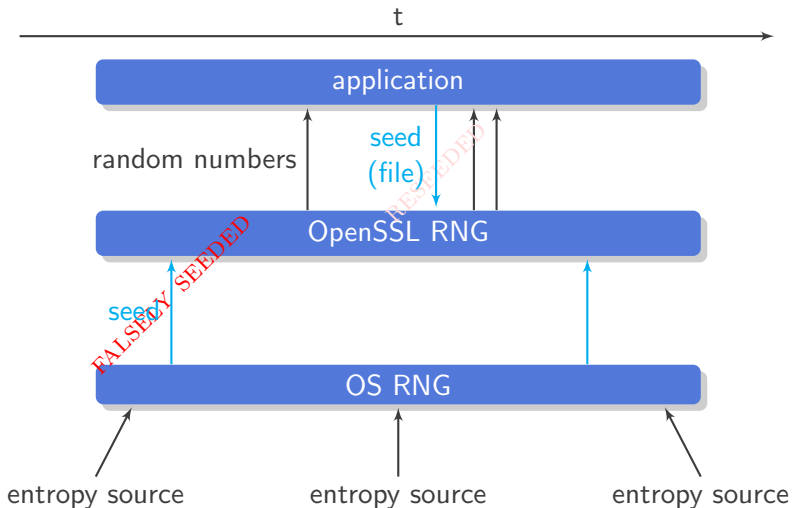




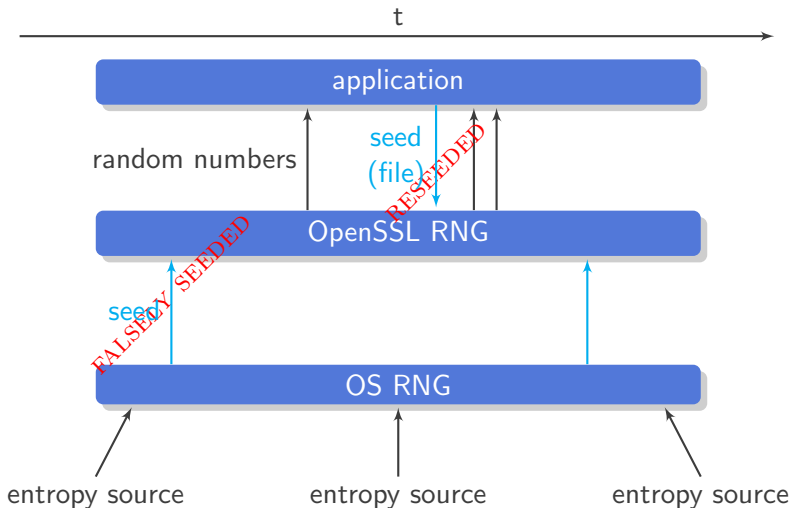
# Reseeded State in Practice



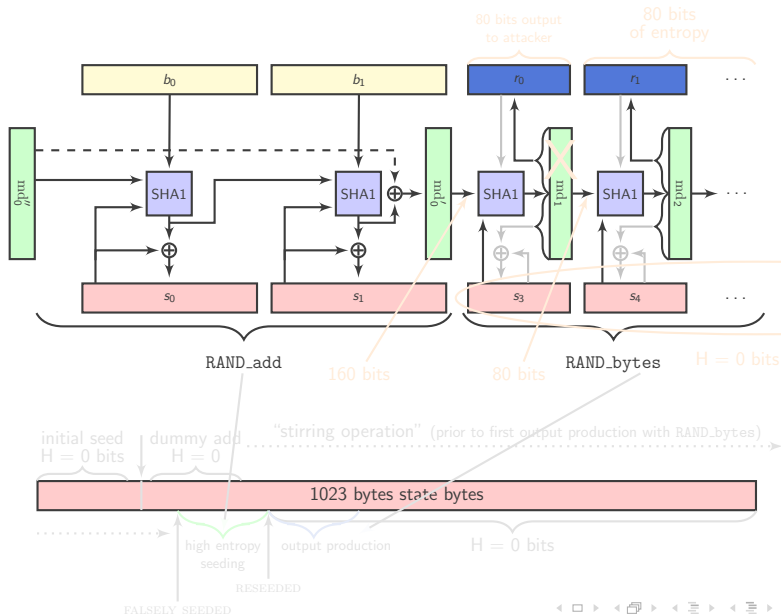
# Reseeded State in Practice



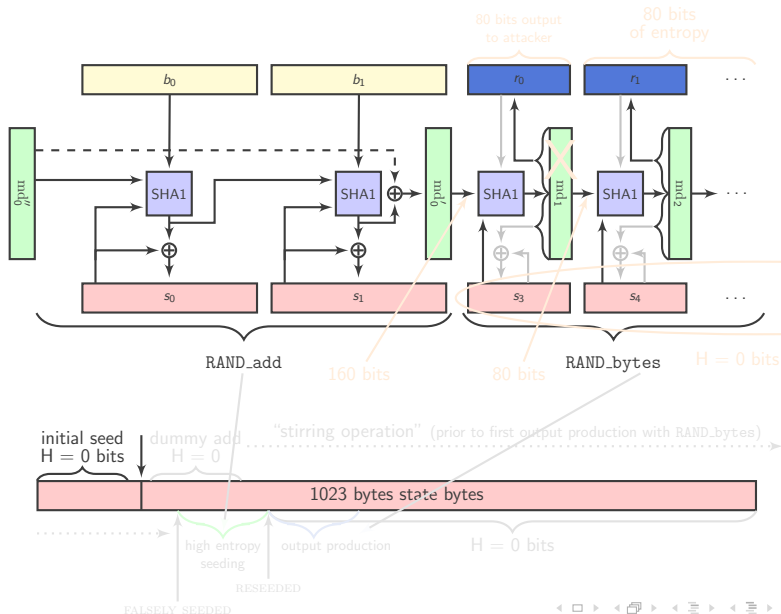
# Reseeded State in Practice



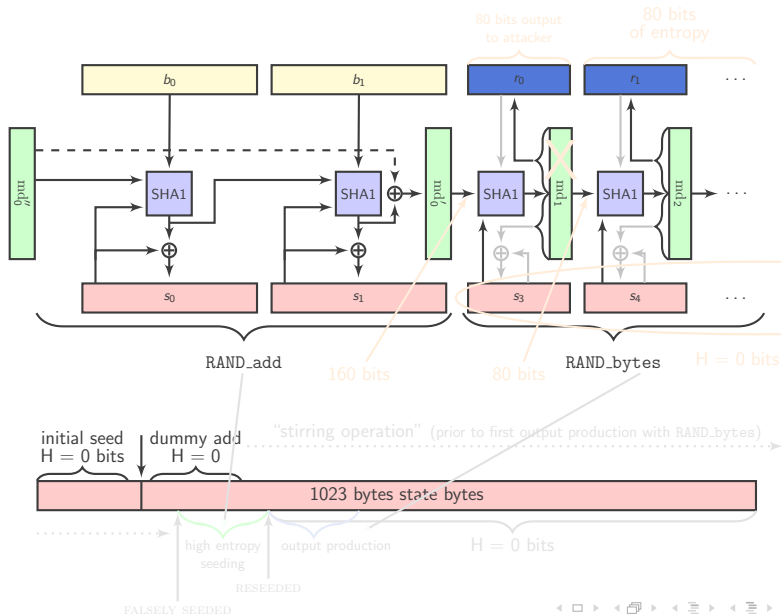
# ELO160 / ELO80 in State RESEEDED



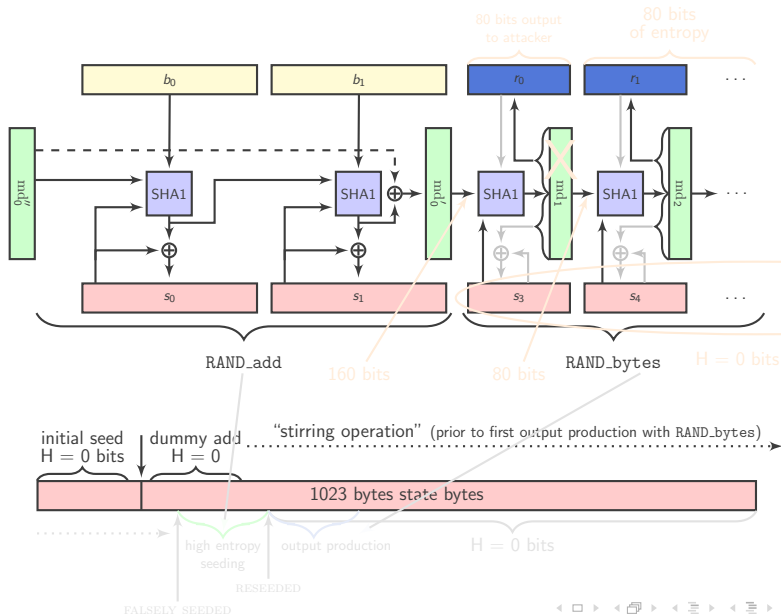
# ELO160 / ELO80 in State RESEEDED



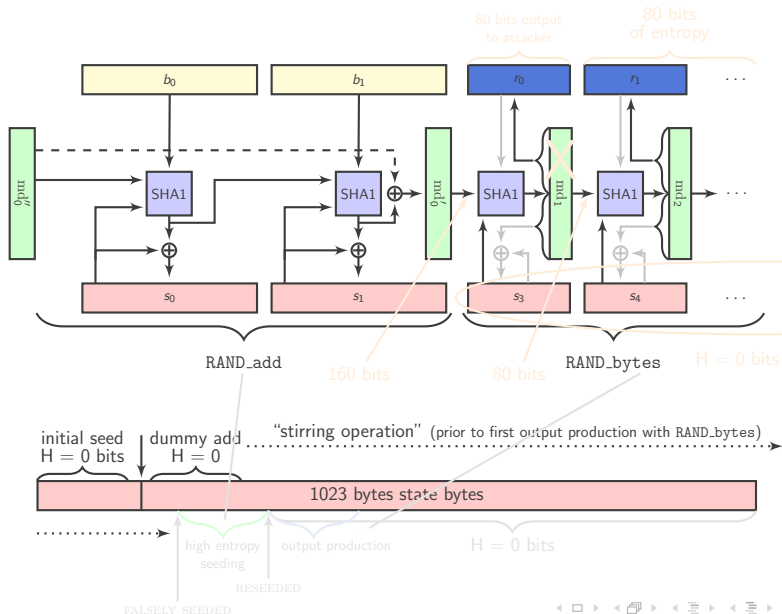
# ELO160 / ELO80 in State RESEEDED



# ELO160 / ELO80 in State RESEEDED

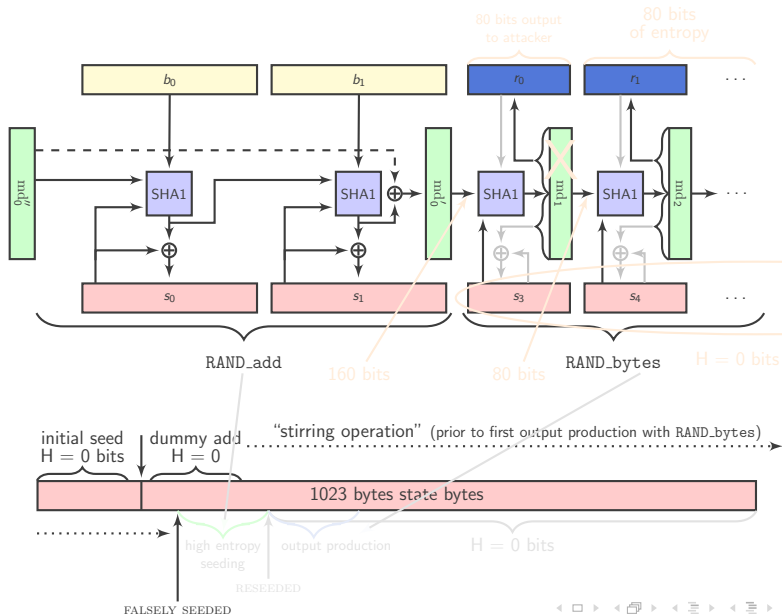


# ELO160 / ELO80 in State RESEEDED

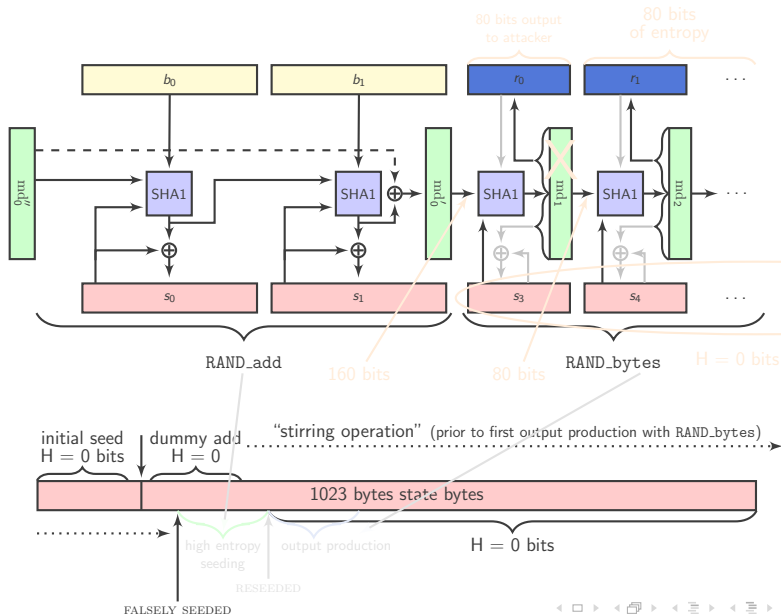




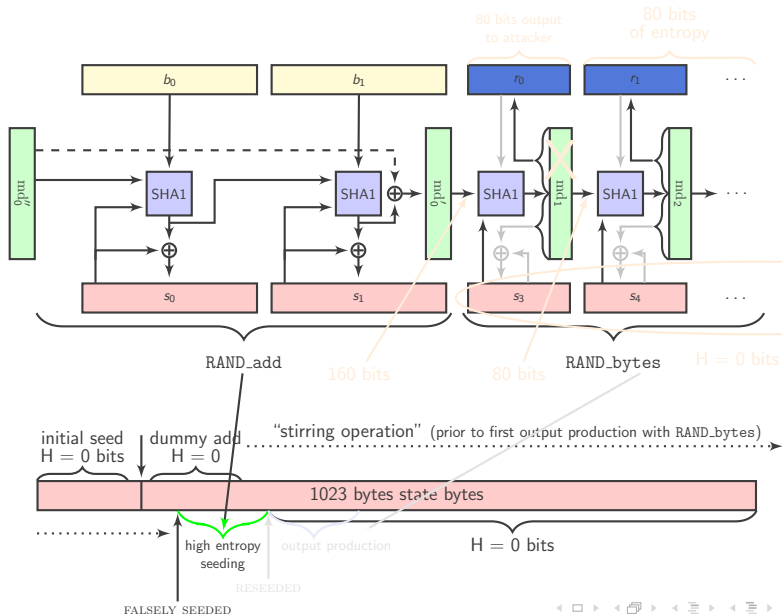
# ELO160 / ELO80 in State RESEEDED



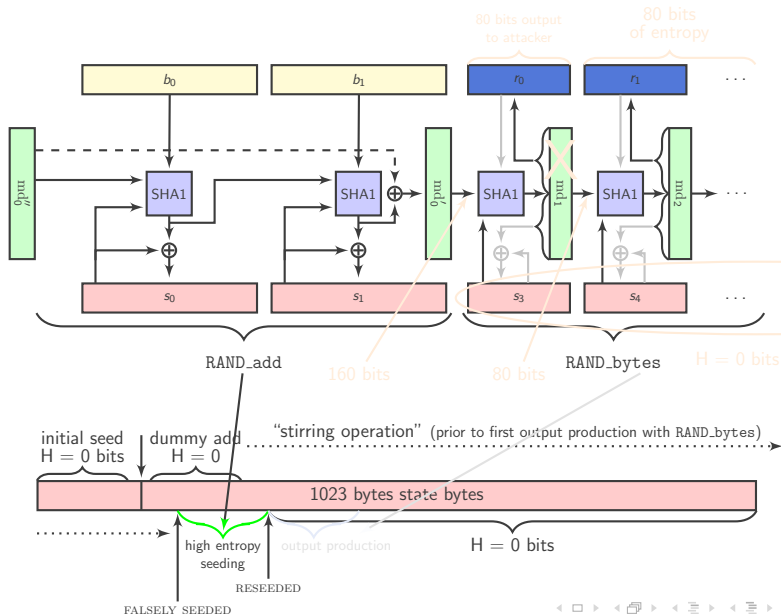
# ELO160 / ELO80 in State RESEEDED



# ELO160 / ELO80 in State RESEEDED

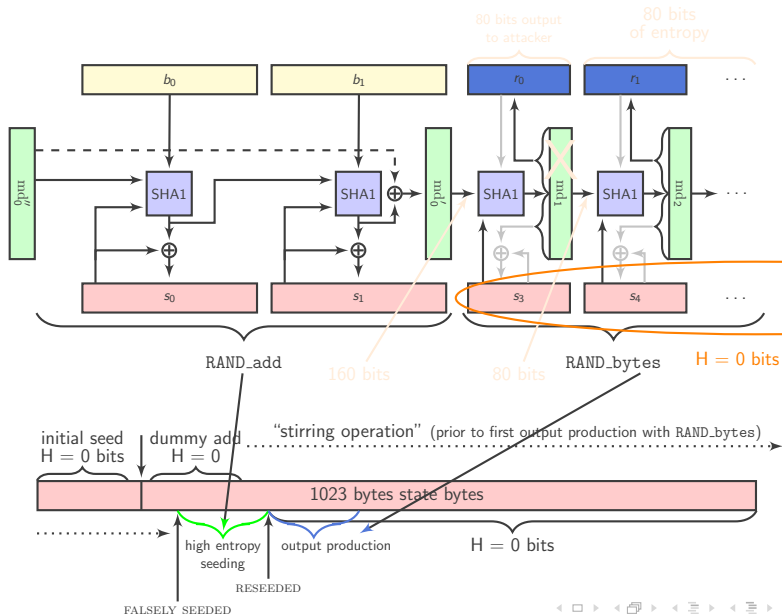


# ELO160 / ELO80 in State RESEEDED

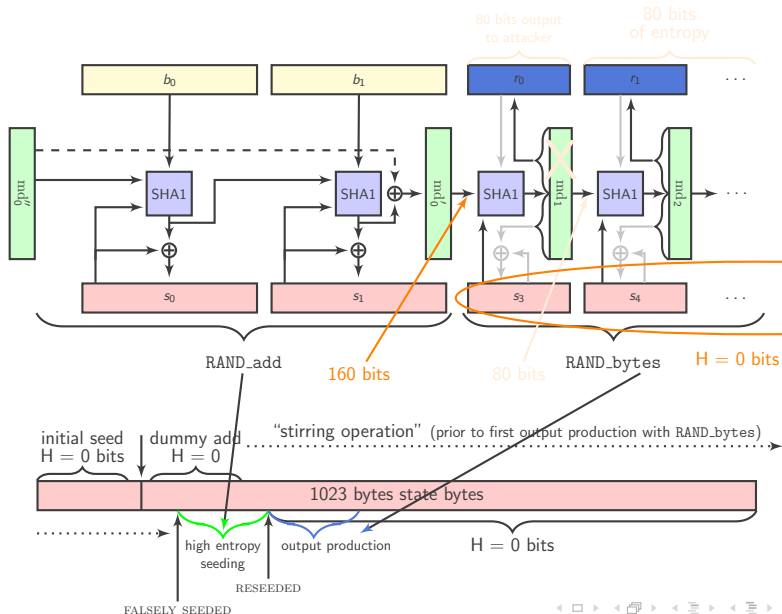




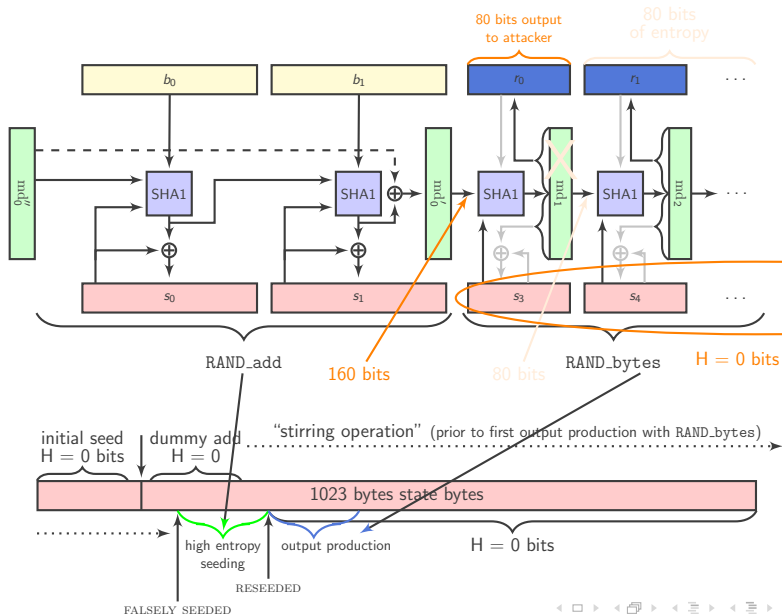
# ELO160 / ELO80 in State RESEEDED



# ELO160 / ELO80 in State RESEEDED

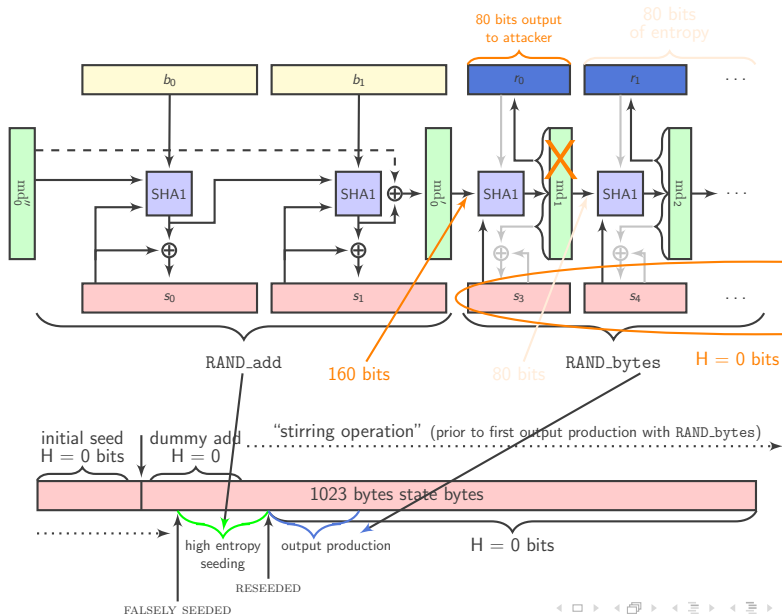


# ELO160 / ELO80 in State RESEEDED

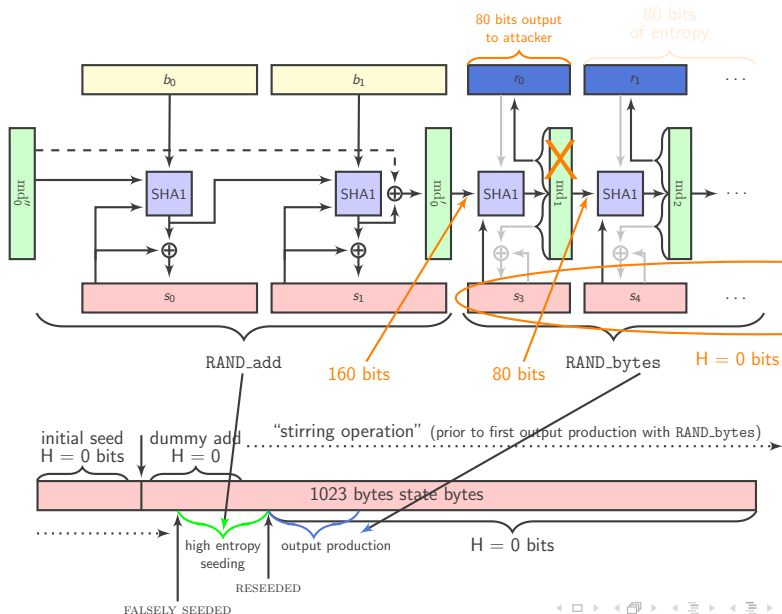




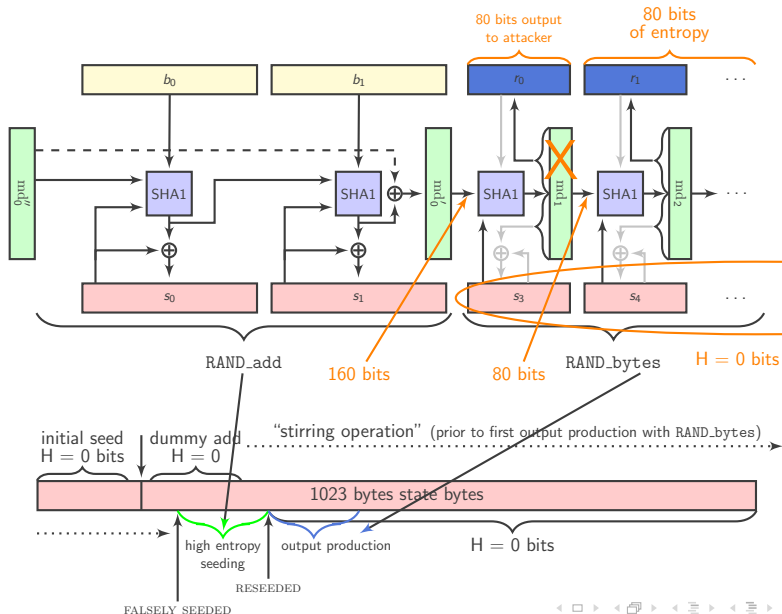
# ELO160 / ELO80 in State RESEEDED



# ELO160 / ELO80 in State RESEEDED



# ELO160 / ELO80 in State RESEEDED



- ELO-240: purely cosmetic
- ELO-160: not exploitable
- ELO-80: only predict output from same call to `RAND_bytes`
- can we do better?  $\approx 2^{80}$  and more realistic conditions?

- ELO-240: purely cosmetic
- ELO-160: not exploitable
- ELO-80: only predict output from same call to `RAND_bytes`
- can we do better?  $\approx 2^{80}$  and more realistic conditions?

- ELO-240: purely cosmetic
- ELO-160: not exploitable
- ELO-80: only predict output from same call to `RAND_bytes`
- can we do better?  $\approx 2^{80}$  and more realistic conditions?


- ELO-240: purely cosmetic
- ELO-160: not exploitable
- ELO-80: only predict output from same call to `RAND_bytes`
- can we do better?  $\approx 2^{80}$  and more realistic conditions?

# State Recovery Attacks




# State Recovery Attack: DEJA-STATE

- RNG in low entropy state
- high entropy reseeding
- in RESEEEDED state
- goal: recover RNG state after reseeding




# State Recovery Attack: DEJA-STATE

- RNG in low entropy state
- high entropy reseeding
- in RESEEEDED state
- goal: recover RNG state after reseeding




# State Recovery Attack: DEJA-STATE

- RNG in low entropy state
- high entropy reseeding
- in RESEEDED state
- goal: recover RNG state after reseeding




# State Recovery Attack: DEJA-STATE

- RNG in low entropy state
- high entropy reseeding
- in RESEDED state
- goal: recover RNG state after reseeding

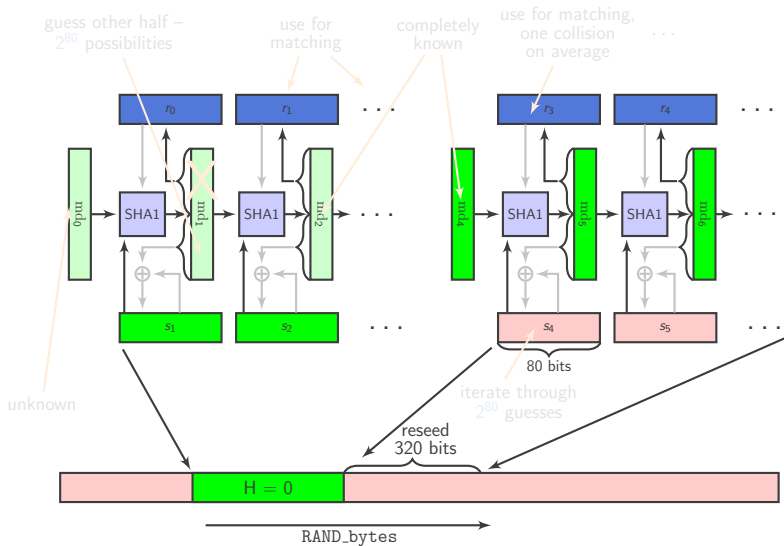


# State Recovery Attack: DEJA-STATE

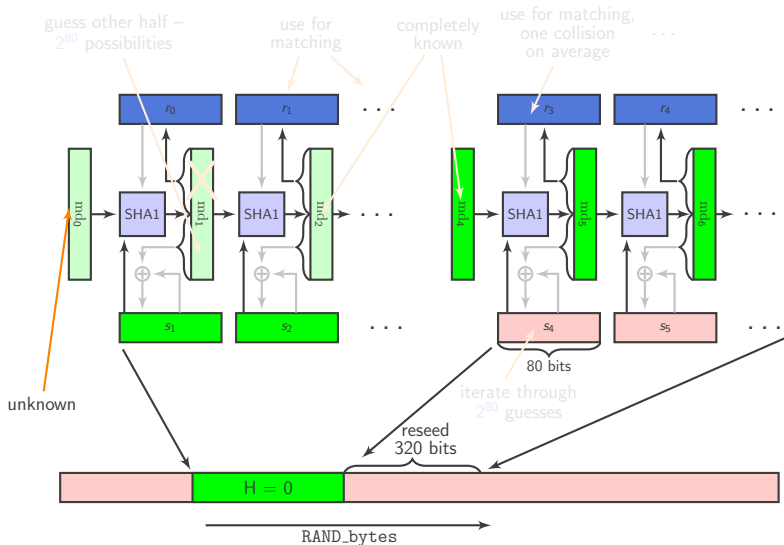
- RNG in low entropy state
- high entropy reseeding
- in RESEEEDED state
- goal: recover RNG state after reseeding



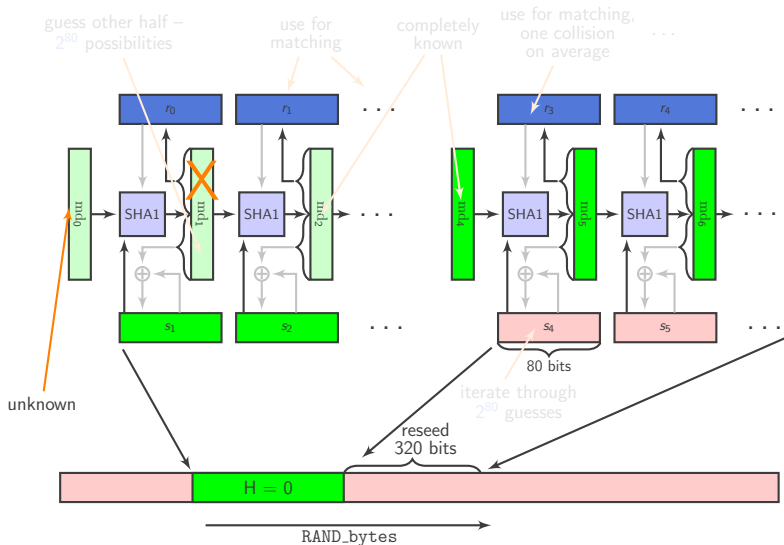
# State Recovery Attack: DEJA-STATE



# State Recovery Attack: DEJA-STATE




# State Recovery Attack: DEJA-STATE

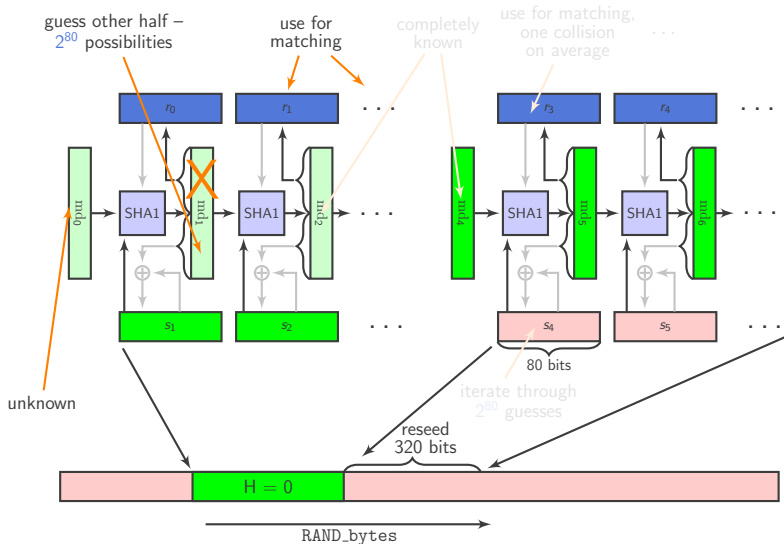




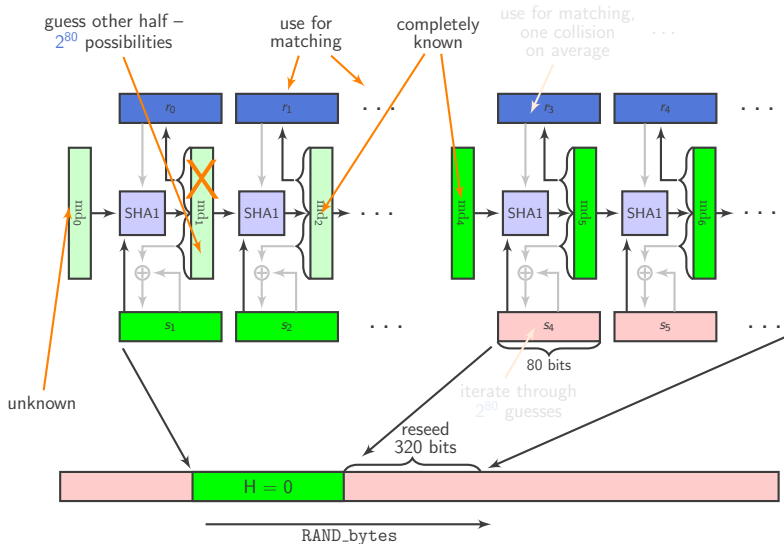
# State Recovery Attack: DEJA-STATE



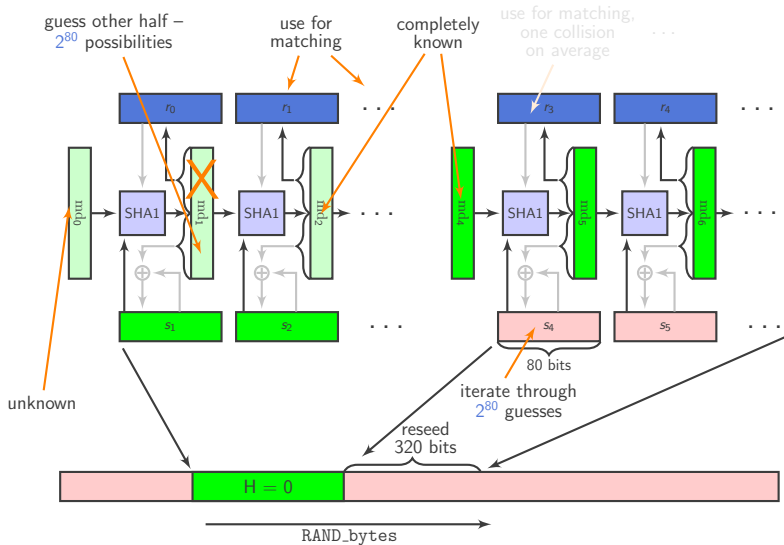
# State Recovery Attack: DEJA-STATE



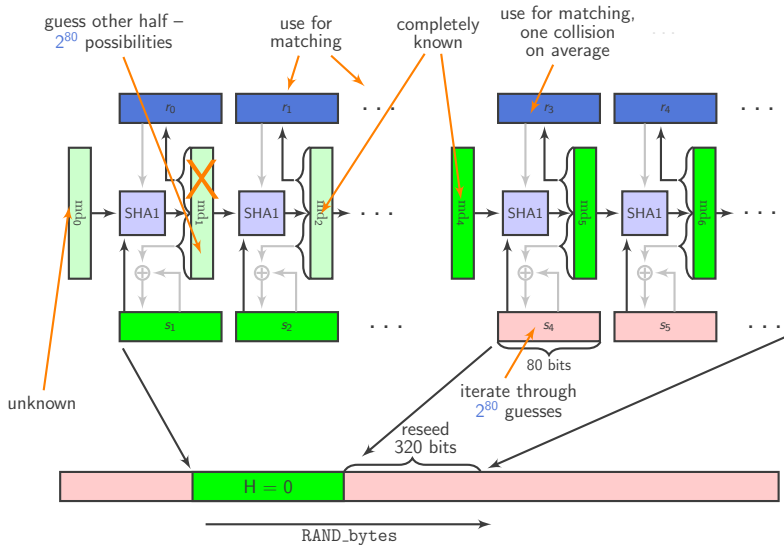
# State Recovery Attack: DEJA-STATE



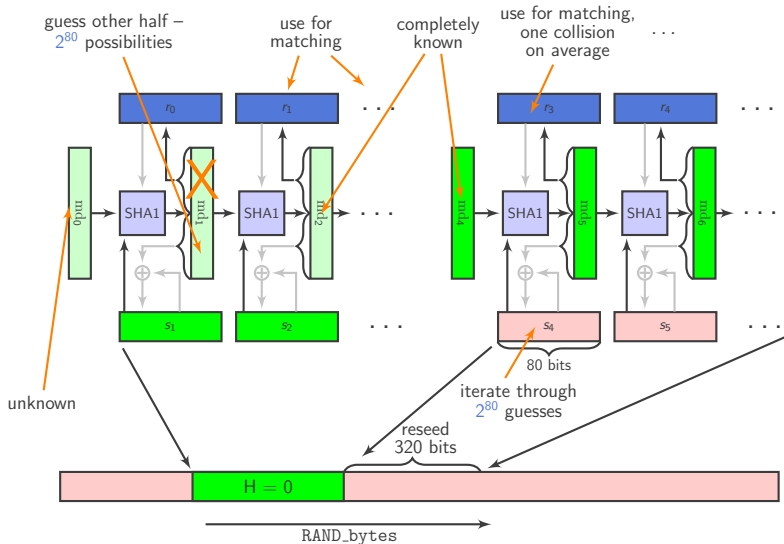
# State Recovery Attack: DEJA-STATE



# State Recovery Attack: DEJA-STATE

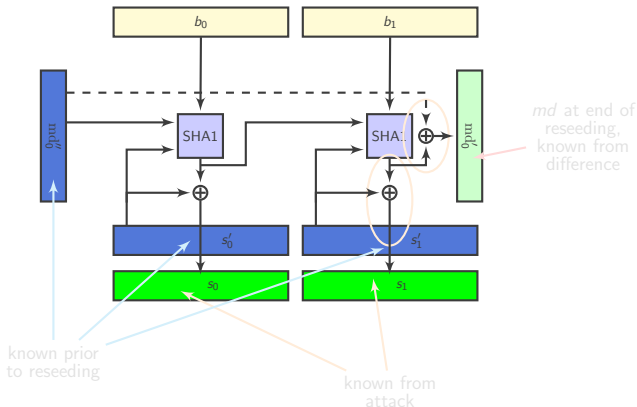


# State Recovery Attack: DEJA-STATE



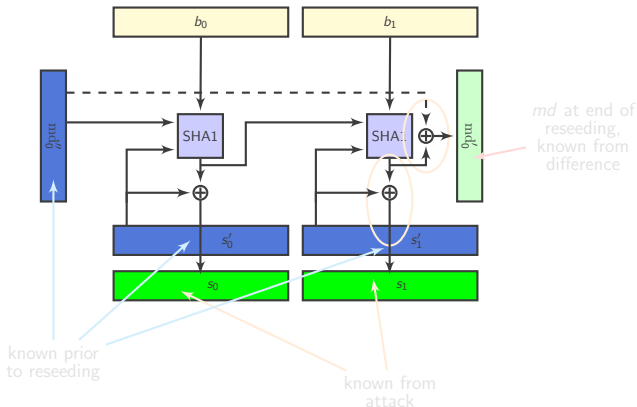
# State Recovery Attack: DEJA-STATE

- state bytes recovered
- now: recovery of  $md$  after the seeding
- revisit the attacked seeding:



# State Recovery Attack: DEJA-STATE

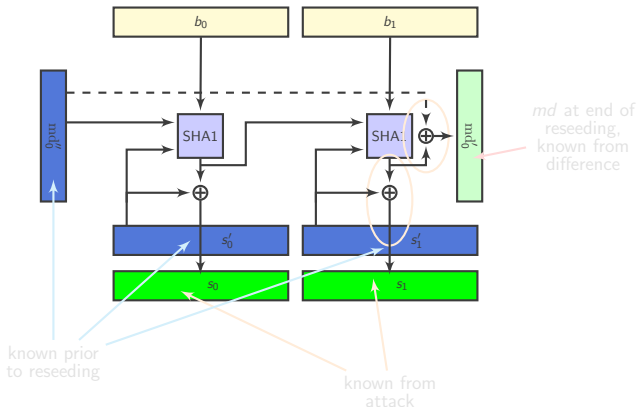
- state bytes recovered
- now: recovery of  $md$  after the seeding
- revisit the attacked seeding:





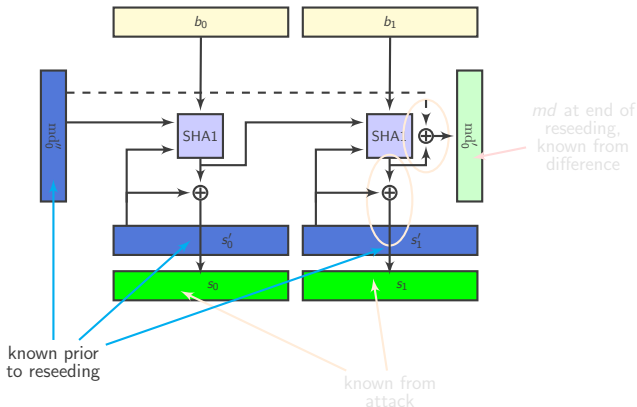
# State Recovery Attack: DEJA-STATE

- state bytes recovered
- now: recovery of  $md$  after the seeding
- revisit the attacked seeding:



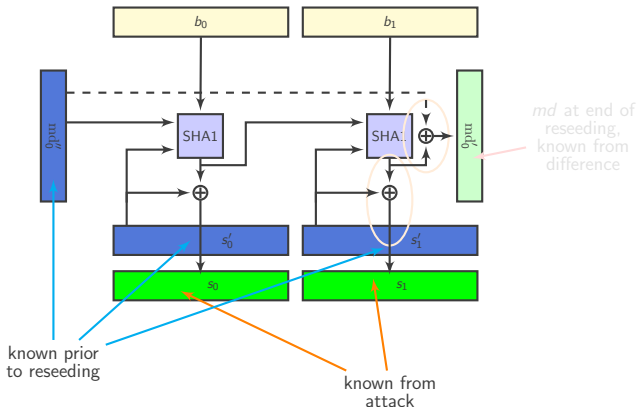
# State Recovery Attack: DEJA-STATE

- state bytes recovered
- now: recovery of  $md$  after the seeding
- revisit the attacked seeding:



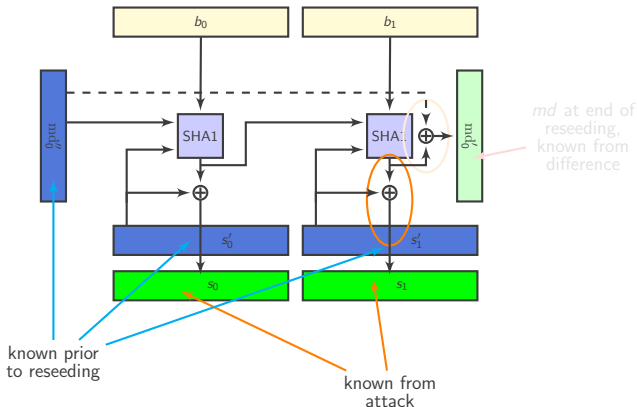
# State Recovery Attack: DEJA-STATE

- state bytes recovered
- now: recovery of  $md$  after the seeding
- revisit the attacked seeding:



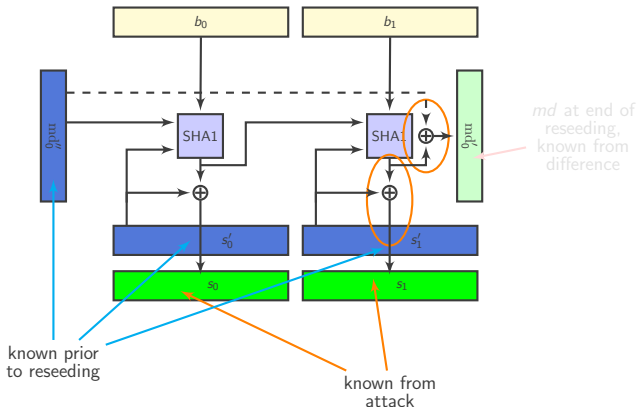
# State Recovery Attack: DEJA-STATE

- state bytes recovered
- now: recovery of  $md$  after the seeding
- revisit the attacked seeding:



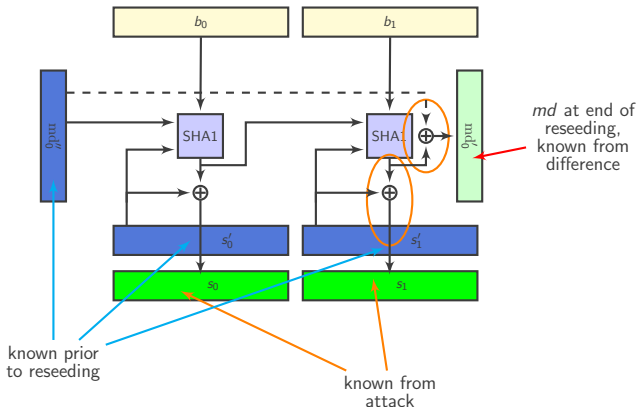
# State Recovery Attack: DEJA-STATE

- state bytes recovered
- now: recovery of  $md$  after the seeding
- revisit the attacked seeding:



# State Recovery Attack: DEJA-STATE

- state bytes recovered
- now: recovery of  $md$  after the seeding
- revisit the attacked seeding:



## Strategies to deal with non-zero initial entropy

- determine state prior to seeding from output
- determine additional entropy during the recovery of md
  - computational effort  $2^{80+x}$

## Strategies to deal with non-zero initial entropy

- determine state prior to seeding from output
- determine additional entropy during the recovery of md
  - computational effort  $2^{80+x}$



## Strategies to deal with non-zero initial entropy

- determine state prior to seeding from output
- determine additional entropy during the recovery of md
  - computational effort  $2^{80+x}$

# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time

# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time

# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time

# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time

# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time

# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time

# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time



# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
    - must be feared to be breakable within a decade (?)
    - will incur considerable costs for a long time

# State Recovery Attack: DEJA-STATE


- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time

# State Recovery Attack: DEJA-STATE

- state after reseeding completely recovered
- condition: attacker receives longer portion of output at specific offset after reseeding
- effort for a 320-bit seed:  $2^{84}$  hash evaluations
- (some tens of bytes in each hash invocation)
- also possible for seed not a multiple of 80 bits
- $2^{80}$  considered “light-weight security”
  - $\approx$  RSA-1024
  - PRESENT light-weight block cipher for RFID applications
  - must be feared to be breakable within a decade (?)
  - will incur considerable costs for a long time

# State Recovery Attack: DEJA-SEED

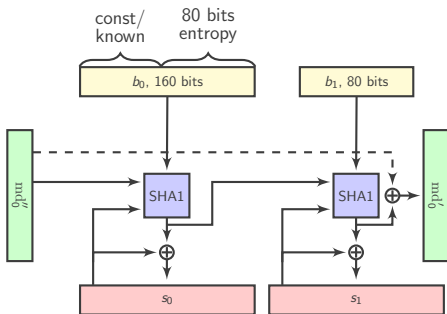
- similar attack, recover also the seed



- synching to  $md$  like in DEJA-STATE
- then iterate through the possible seed values

# State Recovery Attack: DEJA-SEED

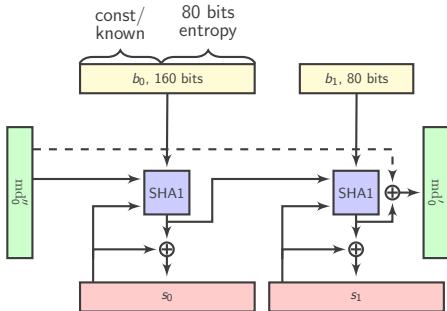
- similar attack, recover also the seed



- synching to *md* like in DEJA-STATE
- then iterate through the possible seed values

# State Recovery Attack: DEJA-SEED

- similar attack, recover also the seed



- synching to  $md$  like in DEJA-STATE
- then iterate through the possible seed values

# Forward Security of Seed Data

- forward security of seed data not a recognized notion
- OpenSSL's RNG: even high entropy seed data potentially recoverable
- block-wise hashing in `RAND_add` is a mistake
- correct: hashing state together with new input
- always inefficient for large RNG states



- forward security of seed data not a recognized notion
- OpenSSL's RNG: even high entropy seed data potentially recoverable
- block-wise hashing in `RAND_add` is a mistake
- correct: hashing state together with new input
- always inefficient for large RNG states

- forward security of seed data not a recognized notion
- OpenSSL's RNG: even high entropy seed data potentially recoverable
- block-wise hashing in `RAND_add` is a mistake
  - correct: hashing state together with new input
  - always inefficient for large RNG states

- forward security of seed data not a recognized notion
- OpenSSL's RNG: even high entropy seed data potentially recoverable
- block-wise hashing in `RAND_add` is a mistake
- correct: hashing state together with new input
- always inefficient for large RNG states

- forward security of seed data not a recognized notion
- OpenSSL's RNG: even high entropy seed data potentially recoverable
- block-wise hashing in `RAND_add` is a mistake
- correct: hashing state together with new input
- always inefficient for large RNG states

# Theoretical Considerations

$$H(f(I, S)) \geq H(S) \text{ and } H(f(I, S)) \geq H(I)$$

- input  $I$ , state  $S$ 
  - RAND\_add fulfills this notion formally
  - but not effectively
  - only useful if whole state is used in output production in a **symmetric way**
  - need definition which considers **entropy of subsequent output** instead of that of the state

$$H(f(I, S)) \geq H(S) \text{ and } H(f(I, S)) \geq H(I)$$

- input  $I$ , state  $S$
- RAND\_add fulfills this notion formally
- but not effectively
- only useful if whole state is used in output production in a **symmetric way**
- need definition which considers **entropy of subsequent output** instead of that of the state

$$H(f(I, S)) \geq H(S) \text{ and } H(f(I, S)) \geq H(I)$$

- input  $I$ , state  $S$
- RAND\_add fulfills this notion formally
- but not effectively
- only useful if whole state is used in output production in a **symmetric way**
- need definition which considers **entropy of subsequent output** instead of that of the state



$$H(f(I, S)) \geq H(S) \text{ and } H(f(I, S)) \geq H(I)$$

- input  $I$ , state  $S$
- RAND\_add fulfills this notion formally
- but not effectively
- only useful if whole state is used in output production in a **symmetric way**
- need definition which considers **entropy of subsequent output** instead of that of the state

$$H(f(I, S)) \geq H(S) \text{ and } H(f(I, S)) \geq H(I)$$

- input  $I$ , state  $S$
- RAND\_add fulfills this notion formally
- but not effectively
- only useful if whole state is used in output production in a **symmetric way**
- need definition which considers **entropy of subsequent output** instead of that of the state

# Formal Vulnerabilities of OpenSSL's RNG

- impaired forward security




- backward security not attempted by RNG itself
  - but when attempted by application, suffers from our attacks



- new notion: forward security of seed data
  - not achieved by OpenSSL's RNG

# Formal Vulnerabilities of OpenSSL's RNG

- impaired forward security




- backward security not attempted by RNG itself
  - but when attempted by application, suffers from our attacks



- new notion: forward security of seed data
  - not achieved by OpenSSL's RNG

# Formal Vulnerabilities of OpenSSL's RNG

- impaired forward security




- backward security not attempted by RNG itself
  - but when attempted by application, suffers from our attacks



- new notion: forward security of seed data
  - not achieved by OpenSSL's RNG

# Formal Vulnerabilities of OpenSSL's RNG

- impaired forward security




- backward security not attempted by RNG itself
  - but when attempted by application, suffers from our attacks




- new notion: forward security of seed data
  - not achieved by OpenSSL's RNG

# Formal Vulnerabilities of OpenSSL's RNG

- impaired forward security




- backward security not attempted by RNG itself
  - but when attempted by application, suffers from our attacks




- new notion: forward security of seed data
  - not achieved by OpenSSL's RNG

# Formal Vulnerabilities of OpenSSL's RNG

- impaired forward security



- backward security not attempted by RNG itself
  - but when attempted by application, suffers from our attacks




- new notion: forward security of seed data
  - not achieved by OpenSSL's RNG




# Formal Vulnerabilities of OpenSSL's RNG

- impaired forward security



- backward security not attempted by RNG itself
  - but when attempted by application, suffers from our attacks



- new notion: forward security of seed data
  - not achieved by OpenSSL's RNG

# Repairing OpenSSL's RNG

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generator
    - as marked in the FIPS version of the library
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
      - as realized in the FIPS version of the library!
    - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse



# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

# Repairing the RNG

- `RAND_pseudo_bytes` must use different state (LESLI)
- use cipher-based generator
  - approved and efficient designs exist
    - e.g. AES / counter mode generators
    - as realized in the FIPS version of the library!
  - more efficient than hash-based, due to hardware support
- ad-hoc repair
  - increase the “entropy flow” beyond 160 bits
  - remove the leakage of half of md
  - forward security of seed-data cannot be efficiently addressed
- so far **no** repair in OpenSSL
- secure wrapper functions (→ paper)
- Note: the forks LibreSSL and BoringSSL are even worse

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested



- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240,ELO160,ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240,ELO160,ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240,ELO160,ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested



- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
    - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

- multiple design errors in OpenSSL's RNG
  - LESLI
  - ELO240, ELO160, ELO80
  - DEJA-STATE, DEJA-SEED
    - effort around  $2^{80}$  hash evaluations
- impact
  - attacks highly application specific
  - relevant for embedded systems
- theoretic insights
  - applicability of the notion of mixing function
  - forward security of seed data
- repairs suggested

Thank you!