



北京大学
PEKING UNIVERSITY

A Practical Cold Boot Attack on RSA Private Keys

Tian Wang¹, Xiaoxin Cui¹, Yewen Ni¹, Dunshan Yu¹, Xiaole Cui², Gang Qu³

¹ Institute of Microelectronics, Peking University, Beijing, China

²Key Lab of Integrated Microsystems, Peking University Shenzhen Graduate School, Shenzhen, China

³ECE Department, University of Maryland, College Park, MD, USA

Email: cuixx@pku.edu.cn ; cuixl@pkusz.edu.cn ; gangqu@umd.edu



Outline

- Cold boot attacks on RSA encryption schemes
 - Cold boot attack and memory decay
 - Existing algorithms to recover RSA private keys
- Practical cold boot attacks on RSA private keys
 - Limitations of the existing algorithms
 - Our improvements
- Evaluation results and discussion
- Conclusion



Outline

- *Cold boot attacks on RSA encryption schemes*
 - Cold boot attack and memory decay
 - Existing algorithms to recover RSA private keys
- Practical cold boot attacks on RSA private keys
 - Limitations of the existing algorithms
 - Our improvements
- Evaluation results and discussion
- Conclusion

Cold boot attack

- J. Halderman et al, USENIX Security Symposium, August 2008
- Memory remanence effect -Data persists after power-off.
- Memory blocks are transferred to other machines to extract the encryption keys .



Fig.1 A typical cold boot attack^[1]

[1] J.A. Halderman et al, “Lest we remember: cold boot attacks on encryption keys,” in *USENIX Security Symposium*, 2008, pp. 45–60



Memory decay




Fig.2 Visualizing memory decay^[1]




Fig.3 Measuring decay for different test systems^[1]

Bit flip!

[1] J.A. Halderman et al, “Lest we remember: cold boot attacks on encryption keys,” in *USENIX Security Symposium*, 2008, pp. 45–60

Memory decay




Fig.2 Visualizing memory decay^[1]




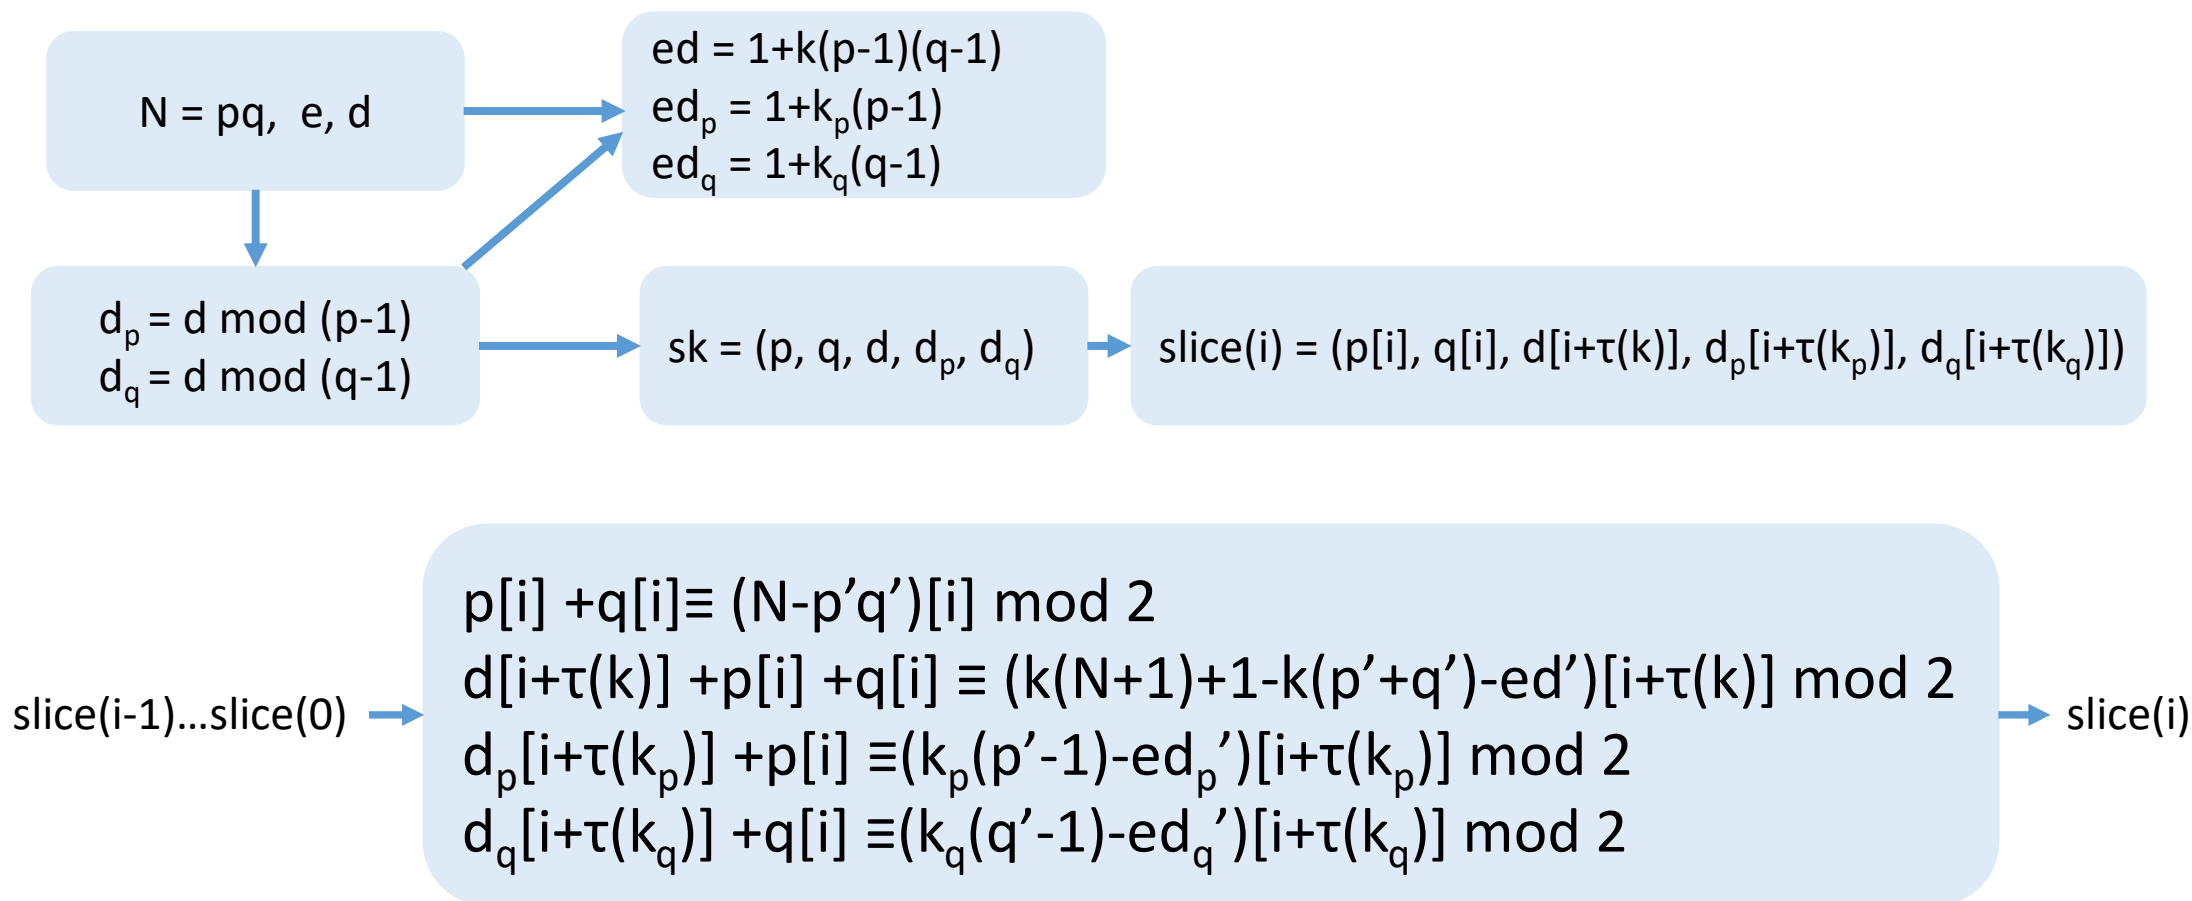
Fig.3 Measuring decay for different test systems^[1]

- Cooling
- Error correction method

[1] J.A. Halderman et al, "Lest we remember: cold boot attacks on encryption keys," in *USENIX Security Symposium*, 2008, pp. 45–60
AsianHOST 2017



Existing algorithms to recover RSA private keys(1/4)





Existing algorithms to recover RSA private keys(2/4)

- Hensel lifting

$$p[i] + q[i] \equiv c_1$$

$$d[i+\tau(k)] + p[i] + q[i] \equiv c_2$$

$$d_p[i+\tau(k_p)] + p[i] \equiv c_3$$

$$d_q[i+\tau(k_q)] + q[i] \equiv c_4$$



Existing algorithms to recover RSA private keys(2/4)

- Hensel lifting

$$p[i] + q[i] \equiv c_1$$

$$d[i + \tau(k)] + p[i] + q[i] \equiv c_2$$

$$d_p[i + \tau(k_p)] + p[i] \equiv c_3$$

$$d_q[i + \tau(k_q)] + q[i] \equiv c_4$$

Slice[0]

$$p[0] = 1$$

$$q[0] = 1$$

$$d[\tau(k)] \dots d[0]$$

$$d_p[\tau(k_p)] \dots d_p[0]$$

$$d_q[\tau(k_q)] \dots d_q[0]$$



Existing algorithms to recover RSA private keys(2/4)

- Hensel lifting

$$p[i] + q[i] \equiv c_1$$

$$d[i + \tau(k)] + p[i] + q[i] \equiv c_2$$

$$d_p[i + \tau(k_p)] + p[i] \equiv c_3$$

$$d_q[i + \tau(k_q)] + q[i] \equiv c_4$$

Slice[1]

$$\begin{aligned} p[1] &= 1 \\ q[1] &= 1 - c_1 \\ d[\tau(k) + 1] &= c_2 - c_1 \\ d_p[\tau(k_p) + 1] &= c_3 - 1 \\ d_q[\tau(k_q) + 1] &= c_4 - q[1] \end{aligned}$$

Slice[0]

$$\begin{aligned} p[0] &= 1 \\ q[0] &= 1 \\ d[\tau(k)] \dots d[0] \\ d_p[\tau(k_p)] \dots d_p[0] \\ d_q[\tau(k_q)] \dots d_q[0] \end{aligned}$$

Existing algorithms to recover RSA private keys(2/4)

- Hensel lifting

$$p[i] + q[i] \equiv c_1$$

$$d[i + \tau(k)] + p[i] + q[i] \equiv c_2$$

$$d_p[i + \tau(k_p)] + p[i] \equiv c_3$$

$$d_q[i + \tau(k_q)] + q[i] \equiv c_4$$

Slice[1]

$$\begin{aligned} p[1] &= 1 \\ q[1] &= 1 - c_1 \\ d[\tau(k) + 1] &= c_2 - c_1 \\ d_p[\tau(k_p) + 1] &= c_3 - 1 \\ d_q[\tau(k_q) + 1] &= c_4 - q[1] \end{aligned}$$

Slice[0]

$$\begin{aligned} p[0] &= 1 \\ q[0] &= 1 \\ d[\tau(k)] \dots d[0] \\ d_p[\tau(k_p)] \dots d_p[0] \\ d_q[\tau(k_q)] \dots d_q[0] \end{aligned}$$

$$\begin{aligned} p[1] &= 0 \\ q[1] &= c_1 \\ d[\tau(k) + 1] &= c_2 - c_1 \\ d_p[\tau(k_p) + 1] &= c_3 \\ d_q[\tau(k_q) + 1] &= c_4 - c_1 \end{aligned}$$



Existing algorithms to recover RSA private keys(2/4)

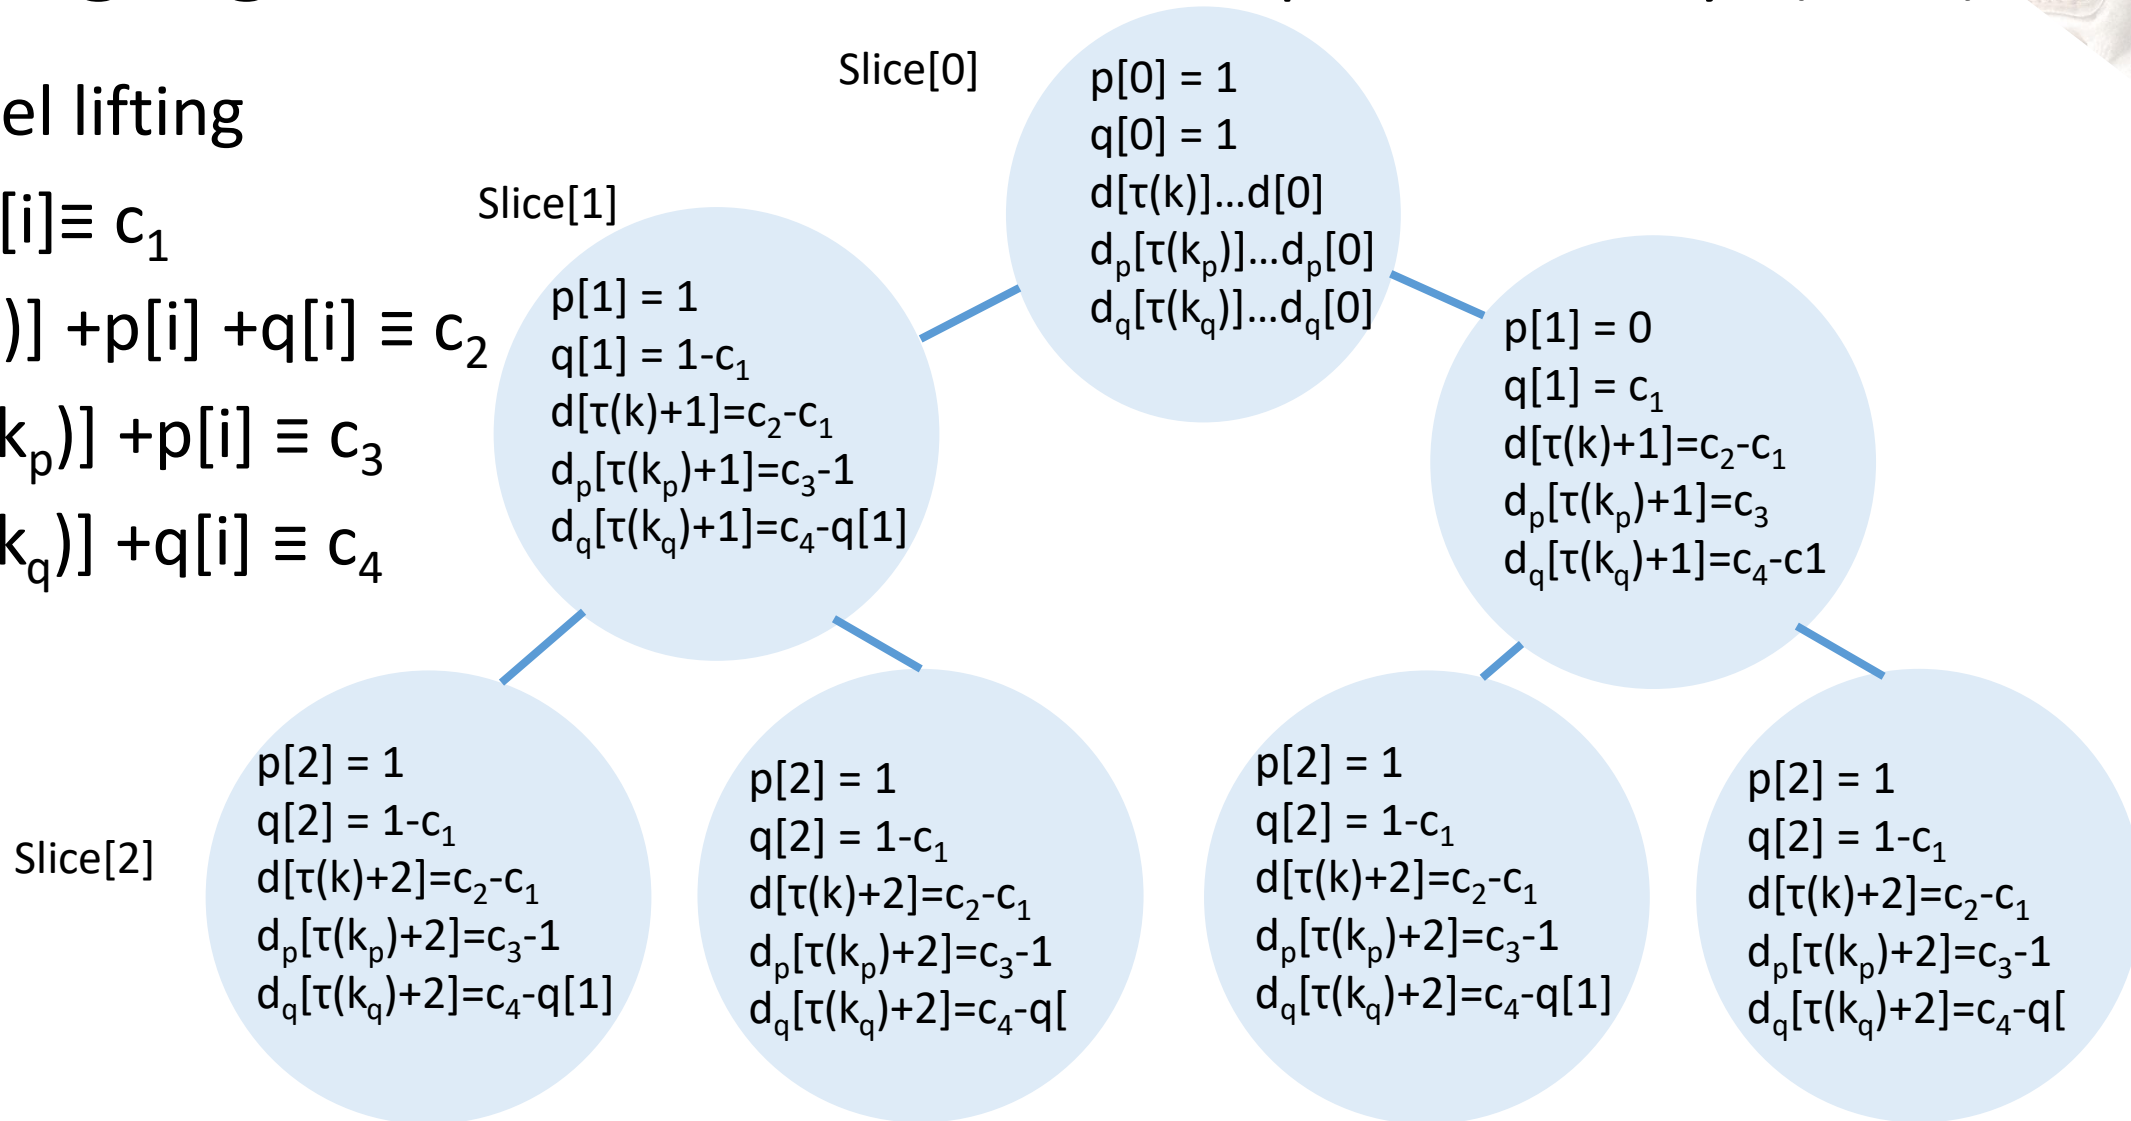
- Hensel lifting


$$p[i] + q[i] \equiv c_1$$

$$d[i + \tau(k)] + p[i] + q[i] \equiv c_2$$

$$d_p[i + \tau(k_p)] + p[i] \equiv c_3$$

$$d_q[i + \tau(k_q)] + q[i] \equiv c_4$$






Existing algorithms to recover RSA private keys(3/4)


HS algorithm^[2]

A fraction $\sigma(\sigma > 0.27)$ of the extracted bits are known.



HMM algorithm^[3]

There exist a flipping probability $\delta(\delta < 0.237)$ for each bit.



[2] N. Heninger and H. Shacham, "Reconstructing RSA private keys from random key bits," *Advances in Cryptology-CRYPTO 2009*. Springer Berlin Heidelberg, 2009, pp.1-17

[3] W. Henecka, A. May and A. Meurer, "Correcting errors in RSA private keys," *Advances in Cryptology-CRYPTO 2010*. Springer Berlin Heidelberg, 2010, pp.351-369.



北京大学
PEKING UNIVERSITY

Existing algorithms to recover RSA private keys(4/4)

HS algorithm

HMM algorithm



北京大学
PEKING UNIVERSITY

Existing algorithms to recover RSA private keys(4/4)

HS algorithm

HMM algorithm

Expansion



Existing algorithms to recover RSA private keys(4/4)

HS algorithm

HMM algorithm

Expansion

Build candidates with Hensel lifting
1 bit a time



Existing algorithms to recover RSA private keys(4/4)

HS algorithm

HMM algorithm

Expansion

Build candidates with Hensel lifting
1 bit a time

Build candidates with Hensel lifting **t** bits a
time



Existing algorithms to recover RSA private keys(4/4)

HS algorithm

Expansion

Build candidates with Hensel lifting
1 bit a time

Pruning

HMM algorithm

Build candidates with Hensel lifting **t** bits a
time



Existing algorithms to recover RSA private keys(4/4)

HS algorithm

Expansion

Build candidates with Hensel lifting
1 bit a time



Pruning

Remove the candidates that don't
match with known key bits

HMM algorithm

Build candidates with Hensel lifting **t** bits a
time



Existing algorithms to recover RSA private keys(4/4)

HS algorithm

Expansion

Build candidates with Hensel lifting
1 bit a time



Pruning

Remove the candidates that don't
match with known key bits

HMM algorithm

Build candidates with Hensel lifting **t** bits a
time



Calculate the number of matching bits



Existing algorithms to recover RSA private keys(4/4)

HS algorithm

Expansion

Build candidates with Hensel lifting
1 bit a time

Pruning

Remove the candidates that don't
match with known key bits

HMM algorithm

Build candidates with Hensel lifting **t** bits a
time

Calculate the number of matching bits

Calculate threshold **C**



Existing algorithms to recover RSA private keys(4/4)

HS algorithm

Expansion

Build candidates with Hensel lifting
1 bit a time

Pruning

Remove the candidates that don't
match with known key bits

HMM algorithm

Build candidates with Hensel lifting **t** bits a
time

Calculate the number of matching bits

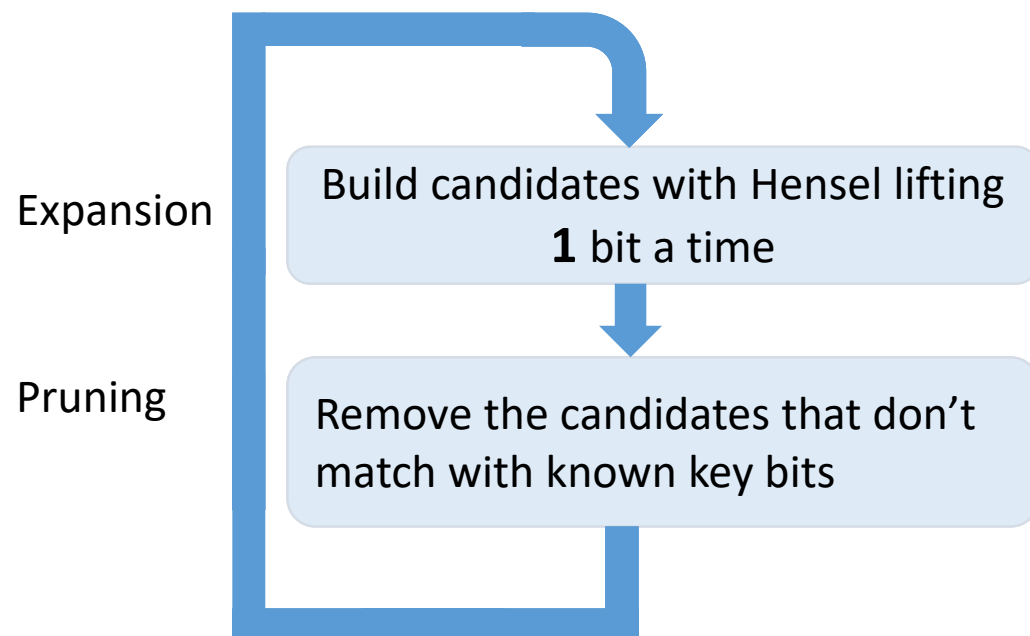
Calculate threshold **C**

Remove the candidates whose number of
matching bits is less than the threshold **C**

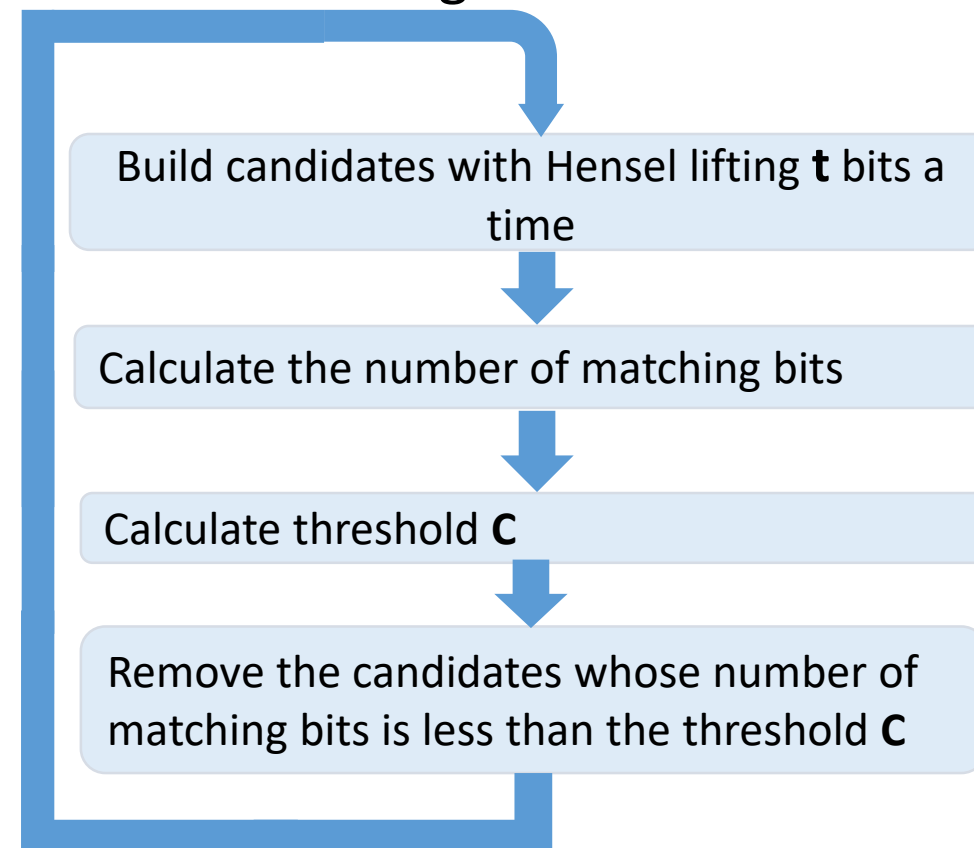


Existing algorithms to recover RSA private keys(4/4)

HS algorithm

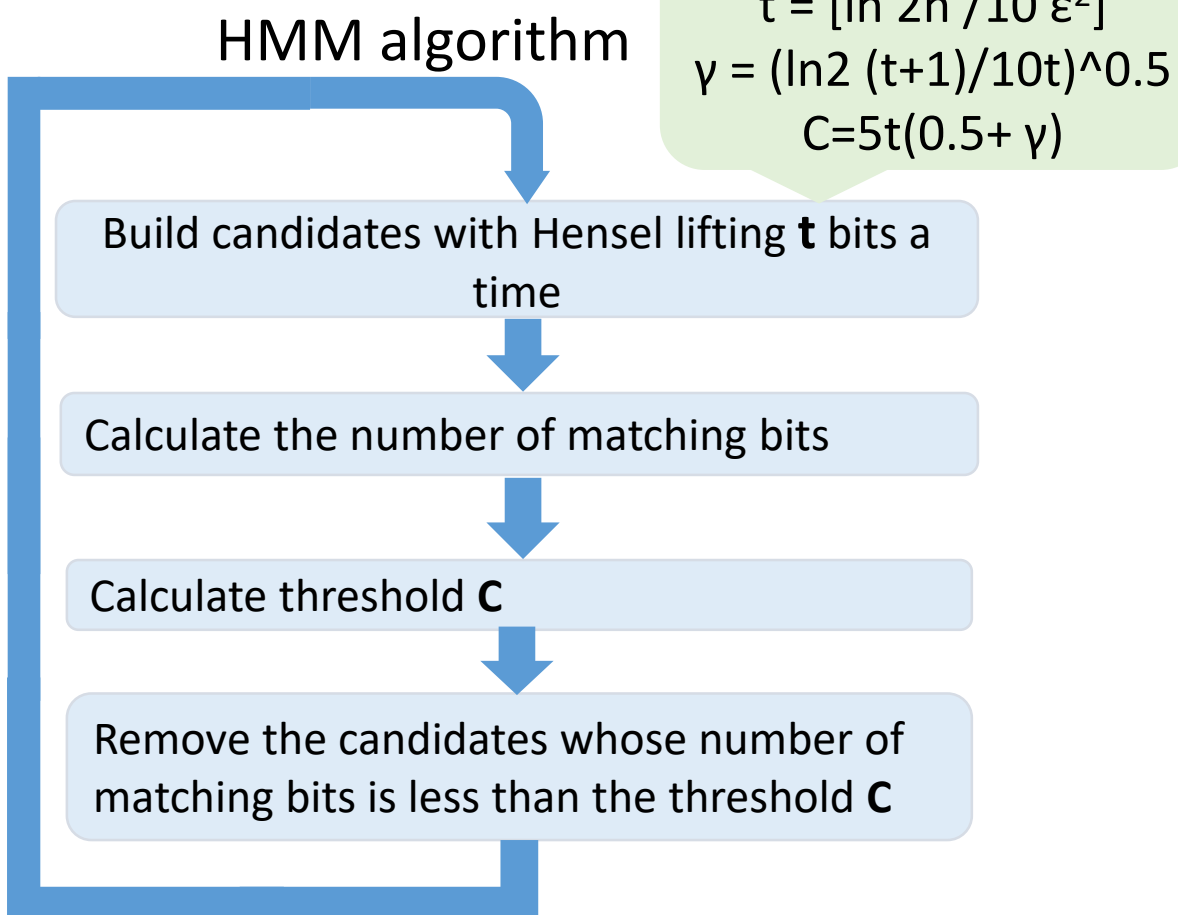
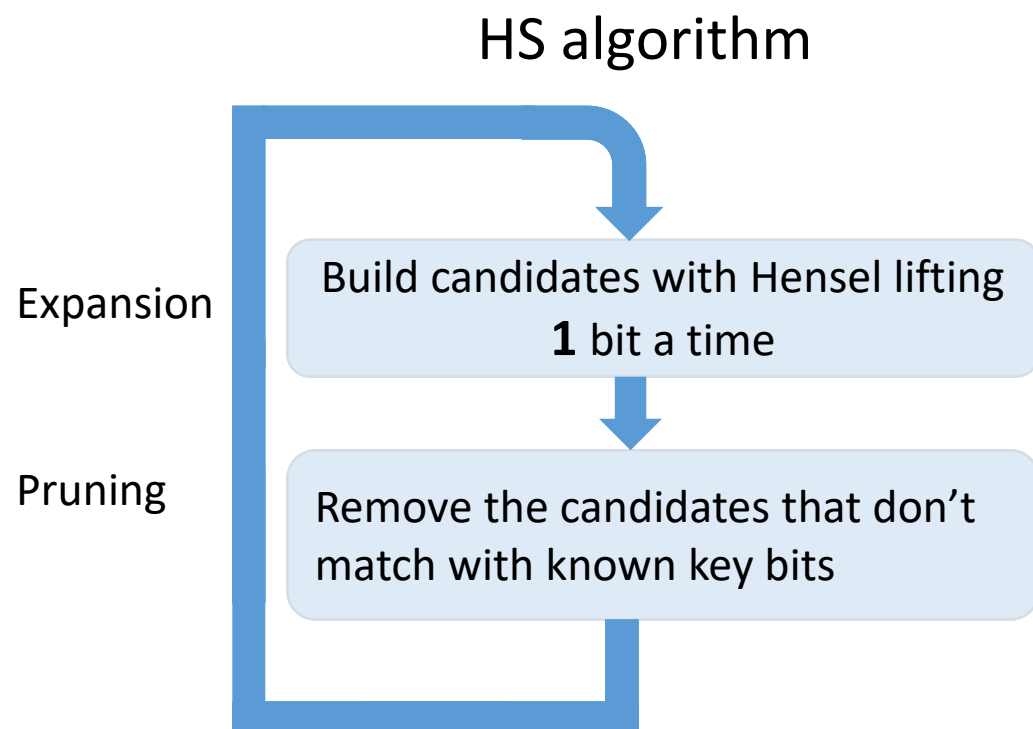


HMM algorithm





Existing algorithms to recover RSA private keys(4/4)





Outline

- Cold boot attacks on RSA encryption schemes
 - Cold boot attack and memory decay
 - Existing algorithms to recover RSA private keys
- *Practical cold boot attacks on RSA private keys*
 - Limitations of the existing algorithms
 - Our improvements
- Evaluation results and discussion
- Conclusion



北京大学
PEKING UNIVERSITY

Limitations of the previous algorithms



Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.



Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.
No bits can be known with absolute certainty.



Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.

No bits can be known with absolute certainty. HS Algorithm



Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.

No bits can be known with absolute certainty. HS Algorithm

- For a given region, the decay is dominated by one direction.

Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.

No bits can be known with absolute certainty. HS Algorithm

- For a given region, the decay is dominated by one direction.

The transition probability in one direction is larger than that of the other direction.

Symmetric decay model is not accurate.

Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.

No bits can be known with absolute certainty. HS Algorithm

- For a given region, the decay is dominated by one direction.

The transition probability in one direction is larger than that of the other direction.

Symmetric decay model is not accurate. HMM Algorithm



Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.

No bits can be known with absolute certainty. HS Algorithm

- For a given region, the decay is dominated by one direction.

The transition probability in one direction is larger than that of the other direction.

Symmetric decay model is not accurate. HMM Algorithm

- The data retention phase is not long enough to mount the attack.



Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.

No bits can be known with absolute certainty. HS Algorithm

- For a given region, the decay is dominated by one direction.

The transition probability in one direction is larger than that of the other direction.

Symmetric decay model is not accurate. HMM Algorithm

- The data retention phase is not long enough to mount the attack.

The flipping probability of the dominating direction can be larger than 0.237.



Limitations of the previous algorithms

- The decay of memory bits from 0 to 1 and from 1 to 0 both exist.

No bits can be known with absolute certainty. HS Algorithm

- For a given region, the decay is dominated by one direction.

The transition probability in one direction is larger than that of the other direction.

Symmetric decay model is not accurate. HMM Algorithm

- The data retention phase is not long enough to mount the attack.

The flipping probability of the dominating direction can be larger than 0.237.

HMM Algorithm



北京大学
PEKING UNIVERSITY

Our improvements based on HMM algorithm(1/4)



北京大学
PEKING UNIVERSITY


Our improvements based on HMM algorithm(1/4)

- Real-life settings.



Our improvements based on HMM algorithm(1/4)


- Real-life settings.





Our improvements based on HMM algorithm(1/4)

- Real-life settings.
- Asymmetric pruning strategy.







Our improvements based on HMM algorithm(1/4)

- Real-life settings.
- Asymmetric pruning strategy.

Our improvements






Our improvements based on HMM algorithm(1/4)

- Real-life settings.
- Asymmetric pruning strategy.

Our improvements



Expansion




Our improvements based on HMM algorithm(1/4)

- Real-life settings.
- Asymmetric pruning strategy.

Our improvements

Build candidates with Hensel lifting t bits a time

Expansion





Our improvements based on HMM algorithm(1/4)


- Real-life settings.
- Asymmetric pruning strategy.

Our improvements

Build candidates with Hensel lifting t bits a time

Expansion


Pruning





Our improvements based on HMM algorithm(1/4)

- Real-life settings.
- Asymmetric pruning strategy.



Our improvements

Build candidates with Hensel lifting t bits a time

Expansion




Calculate the number of matching bit 0 and matching bit 1 respectively

Pruning



Our improvements based on HMM algorithm(1/4)

- Real-life settings.
- Asymmetric pruning strategy.



Our improvements

Build candidates with Hensel lifting t bits a time

Expansion

Calculate the number of matching bit 0 and matching bit 1 respectively


Pruning

Calculate two thresholds for bit 0 and bit 1 C_0 and C_1



Our improvements based on HMM algorithm(1/4)

- Real-life settings.
- Asymmetric pruning strategy.



Our improvements

Build candidates with Hensel lifting t bits a time

Expansion

Calculate the number of matching bit 0 and matching bit 1 respectively

Pruning


Calculate two thresholds for bit 0 and bit 1 C_0 and C_1

Remove the candidates whose number of matching bit 0 (or 1) is less than the corresponding threshold $C_0(C_1)$



Our improvements based on HMM algorithm(1/4)

- Real-life settings.
- Asymmetric pruning strategy.



Our improvements

Build candidates with Hensel lifting t bits a time


Expansion

Calculate the number of matching bit 0 and matching bit 1 respectively

Pruning

Calculate two thresholds for bit 0 and bit 1 C_0 and C_1

Remove the candidates whose number of matching bit 0 (or 1) is less than the corresponding threshold $C_0(C_1)$



Our improvements based on HMM algorithm(2/4)

The Improved HMM Algorithm	Pruning:
Input: (N,e), erroneous sk, error probabilities α and β	For $i = 1$ to 2^t
Output: the correct sk.	--Count the number of matching bit 0(1) $X_{i0(1)}$ of gc_i in slice((round-1)t+1)... slice(round t) with the given erroneous sk.
Initialization: Calculate (k, k_p, k_q) and slice(0)	--Compute the threshold $C_{i0(1)}$ in the ith round. If $(X_{i0} \leq C_{i0} \text{ or } X_{i1} \leq C_{i1})$ Discard the candidate. Else Keep the candidate
For round = 1 to $n/(2t)+1$	end for
Expansion: Expand t times for every candidate with slice(0)...slice(round-1)t using the Hensel lifting, which generates 2^t different partial candidates $gc_1, gc_2, \dots, gc_{2^t}$ in slice((round-1)t+1)... slice(round t).	end for
	Finalization: Run trial decryption to determine the correct private key.

Our improvements based on HMM algorithm(3/4)

- Notation

For any $\varepsilon_0 \varepsilon_1 > 0$, (N, e) as the n -bit modulus and public key.

Given flipping probability $\alpha, \beta, \eta = \min((1 - \alpha) \varepsilon_1^2, \varepsilon_0^2)$

$p_{1(0)}$: the error probability of bit 1(0) in the extracted key

$N_{0(1)}$: the number of bit 0 (1) in each t slices of the extracted key

$N_{0(1)}'$: the number of bit 0 (1) in each t slices of the candidate key

- Choice of parameters

We choose t, t_0, t_1 so that $t = \lceil \ln 2n / 5\eta \rceil$ $t = 1 + t_0 + t_1$

Our improvements based on HMM algorithm(4/4)

- Performance analysis

$$\gamma_1 = (\ln 2 (t_1+1)/2N_1)^{0.5} \quad \gamma_0 = (\ln 2 (t_0+1)/2N_0')^{0.5}$$

$$C_1 = N_1(0.5 + \gamma_1) \quad C_0 = N_0'(N_0/5t + \gamma_0)$$

$$\text{If } p_1 \leq 0.5 - \gamma_1 - \varepsilon_1, \beta \leq 1 - N_0' - \gamma_0 - \varepsilon_0$$

The algorithm can output the correct sk in polynomial time $O(n^{2+\ln 2/2.5\eta})$ with success probability greater than $1-5\eta/2\ln(2n)-1/n$.

*success probability: The possibility that correct candidate isn't pruned.



Outline

- Cold boot attacks on RSA encryption schemes
 - Cold boot attack and memory decay
 - Existing algorithms to recover RSA private keys
- Practical cold boot attacks on RSA private keys
 - Limitations of the existing algorithms
 - Our improvements
- *Evaluation results and discussion*
- Conclusion

Evaluation results(1/3)

- $e = 2^{16}+1$
- $0.05 \leq \alpha \leq 0.4$ $\beta = 0.001, 0.01$
- 100 different 1024-bit private keys
- 100 repetitions for each private key


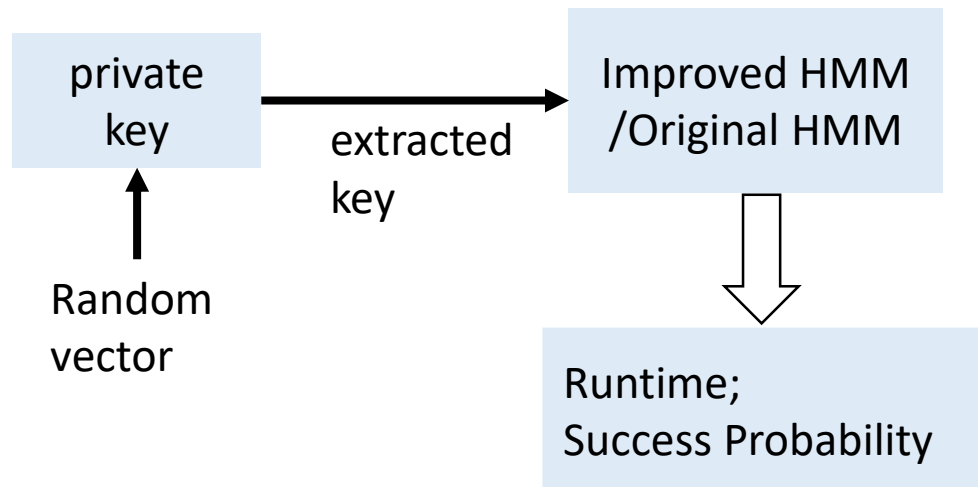


Fig. 4 The success probability and runtime for different combinations of (t_0, t_1) when $\alpha = 0.25$, $\beta = 0.001$ and $t = 11$.



Evaluation results(2/3)





Fig. 5 Success probability and runtime of the improved HMM algorithm(a) and the original HMM algorithm(b)

Optimized t!

Evaluation results(3/3)




Fig. 5 Success probability and runtime of the improved HMM algorithm and the original HMM algorithm

- $\alpha \uparrow$ or $\beta \uparrow \rightarrow t \uparrow \rightarrow$ success probability \downarrow
- $\alpha < 0.2$
- $\alpha > 0.2$

Discussion

- $C = C_1' + C_0' = N_0' (0.5 + \gamma) + (5t - N_0') (0.5 + \gamma)$

- $C_1 = N_1 (0.5 + \gamma_1) \quad C_0 = N_0' (N_0 / 5t + \gamma_0)$

- $\alpha < 0.2 \quad C_1' \approx C_1 \quad C_0' < C_0$

$$P(X_{c0} + X_{c1} > C_0' + C_1') > P(X_{c0} + X_{c1} > C_0 + C_1)$$

$$> P(X_{c0} > C_0 \cap X_{c1} > C_1)$$

- $\alpha > 0.2 \quad C_1' > C_1$

$$P(X_{c0} + X_{c1} > C_0' + C_1') < P(X_{c0} > C_0 \cap X_{c1} > C_1)$$




Fig. 6 Estimated threshold parameters of the improved and original HMM algorithm for different error probabilities α when $\beta = 0.001$



Outline

- Cold boot attacks on RSA encryption schemes
 - Cold boot attack and memory decay
 - Existing algorithms to recover RSA private keys
- Practical cold boot attacks on RSA private keys
 - Limitations of the existing algorithms
 - Our improvements
- Evaluation results and discussion
- *Conclusion*



Conclusion

- A practical improvement of the cold boot attack based on HMM algorithm
 - Real life settings
 - Asymmetric pruning strategy
- Performance analysis



北京大学
PEKING UNIVERSITY



Thank you!