Public Keys

Arjen K. Lenstra (EPFL, Switzerland)
James P. Hughes (Self, Palo Alto, USA)
Maxime Augier (EPFL, Switzerland)
Joppe W. Bos (EPFL, Switzerland)
Thorsten Kleinjung (EPFL, Switzerland)
Christophe Wachter (EPFL, Switzerland)

Insert clip from RSA Cryptographer's panel

http://www.youtube.com/watch?v=y5FeJ6DEaJw

Agenda

- What was collected (and not collected)
- What was computed
- Results
- Discussion
- Conclusion

What we collected

- Openly accessible public keys repositories
  - Static keys (no sniffing, crawling, etc.)
  - MIT PGP Public Key Server
  - EFF SSL Observatory
  - Other keys
- 11.7 million public keys contains
  - 6.4 million distinct RSA moduli.
  - 3.2 million ElGamal keys
  - 3.2 million DSA keys
  - One ECDSA key
- Debian OpenSSL vulnerability were discarded

Results: Duplicates

- Owners may breach each other's security.
- ElGamal and DSA keys
  - a few duplicates with seemingly unrelated owners.
- RSA
  - 6.6 million distinct X.509 moduli
    - certificates and PGP keys
  - 270 thousand (4%) share their RSA modulus.
    - Same moduli used from 2 to 16k times, average 4.
    - Many duplicates occur because of resigning
  - Some duplicates seem to *not* be related
- One PGP duplicate was verified not related

What we computed

- Calculate the GCD of distinct moduli
  - If composite, backtrack
  - If prime, recovered factor
  - If 1 continue
  - Multiply together ensuring no squares
- Implementation
  - The GNU Multiple Precision Arithmetic Library
  - Low memory requirements
- Effort is Subquadratic
- Final integers
  - 10M Moduli
  - 2^30 bytes in length (1GB)
  - 2-3 hours on a Macbook

Trivial example

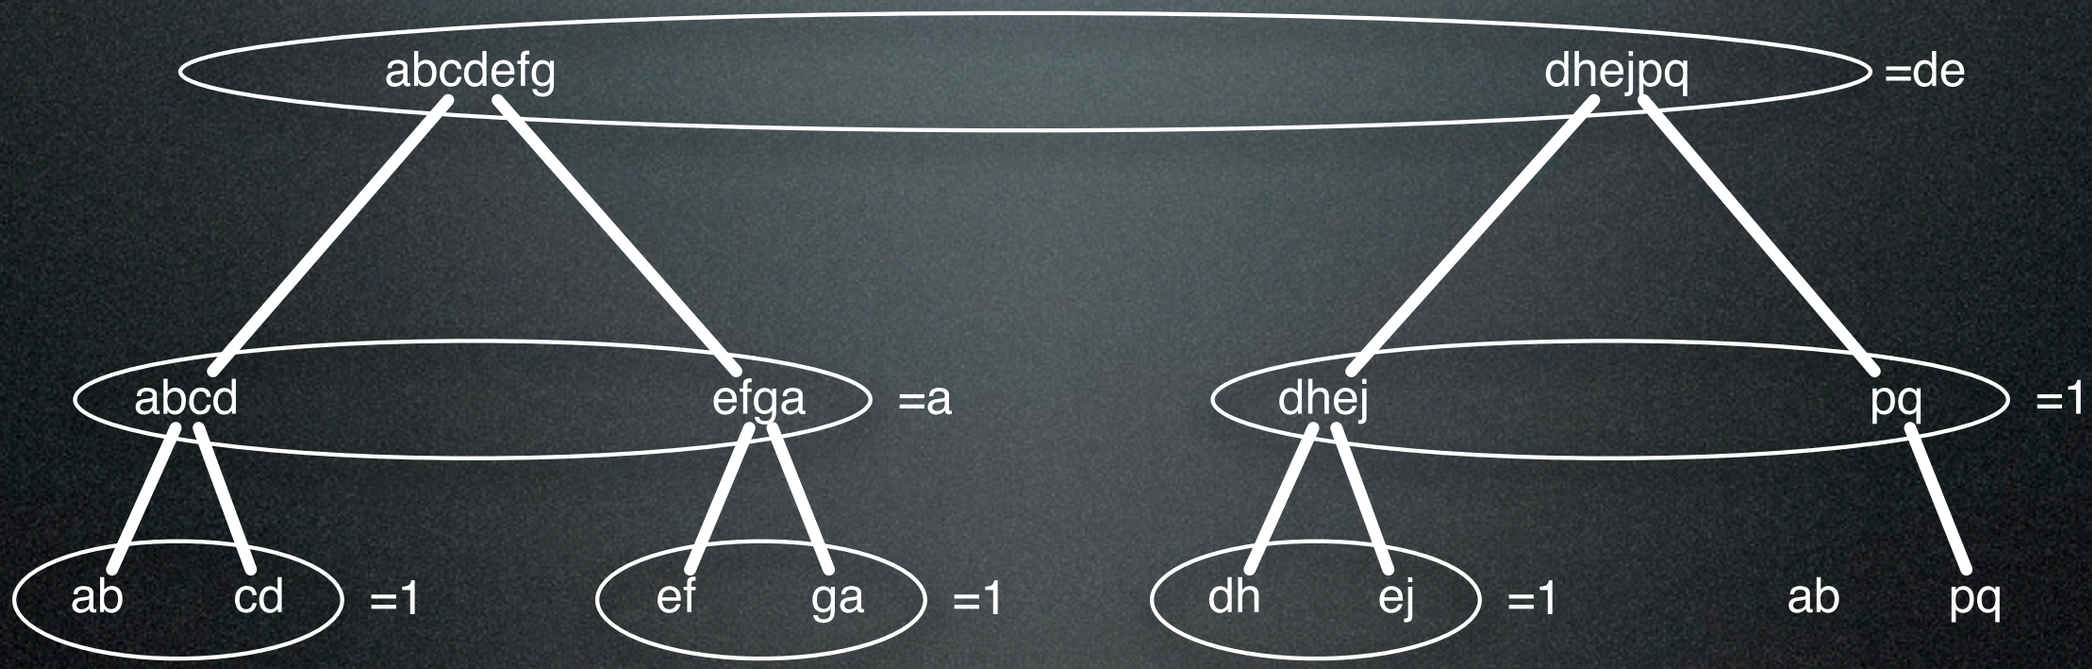

Resulting graph

- Factors = nodes
- Moduli = edge
  - Number = duplicates
- Discard secure keys

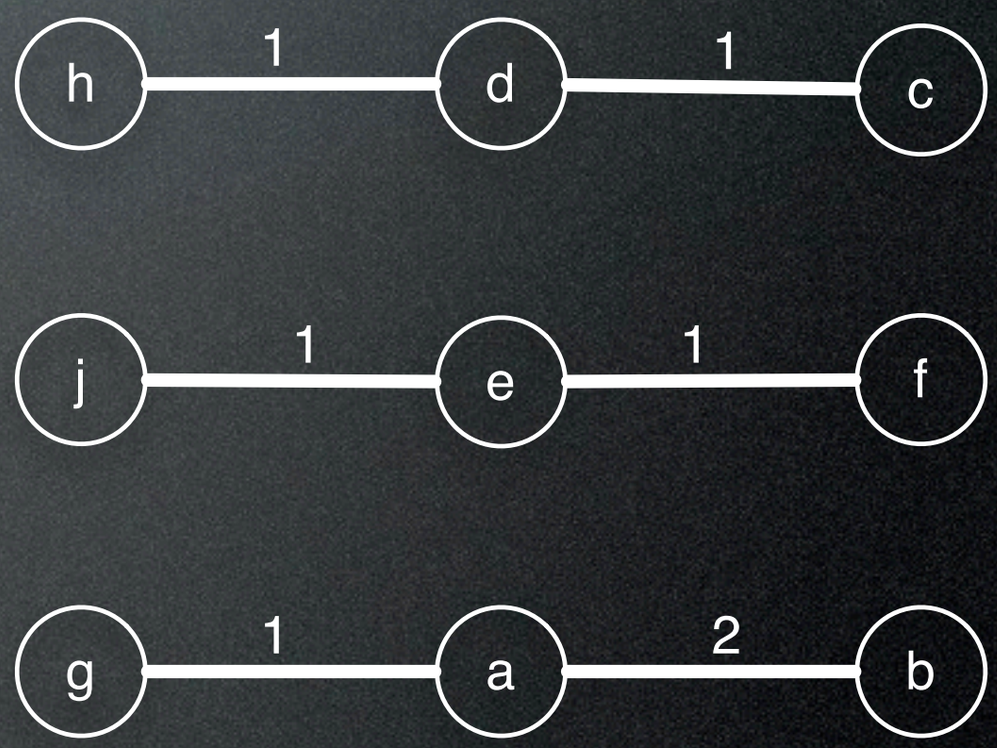

Resulting graph

- Factors = nodes
- Moduli = edge
  - Number = duplicates
- Discard secure keys
- Example
  - three clusters

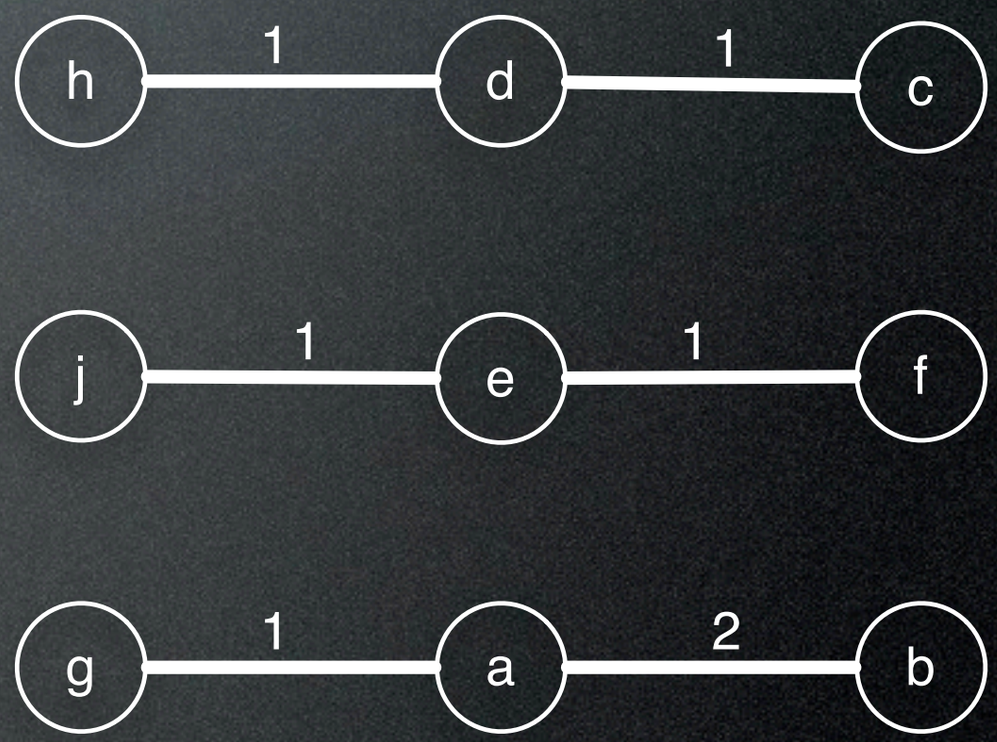

Results: Recoverable keys

- Factors of 12,934 moduli of 1024 bits were recovered
  - 5,250 moduli use SHA1 and not expired
- Factors of 10 moduli of 2048 bits were recovered

- Early conclusions
  - Multiple Vendors
  - Each cluster was the same vendor
  - None of the keys from common eCommerce sites

- Multiple Causes
  - First prime
  - K9
  - Chain

Most common failure

- First prime common
  - Some entropy in second prime
- Initialization from common seed
  - Heninger, et al., “Mining Your Ps and Qs” (2012)

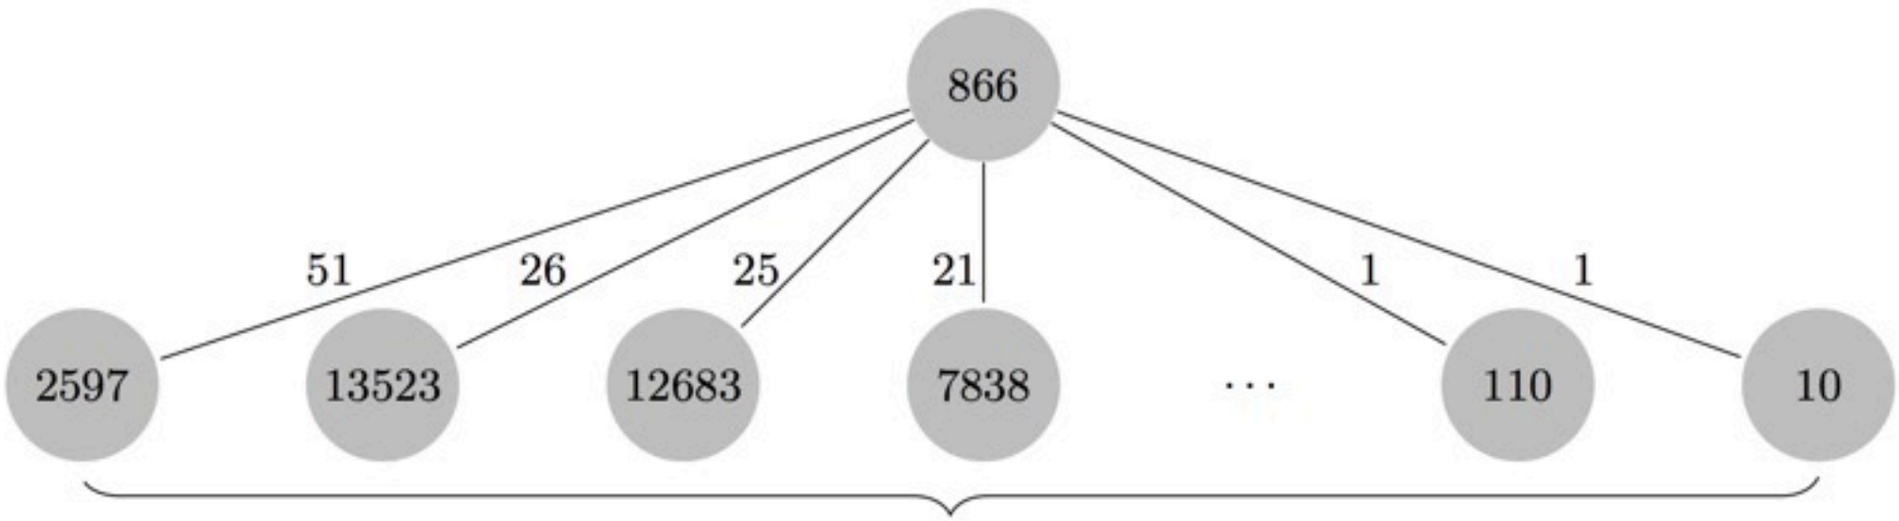

K9: 687 keys from 9 primes

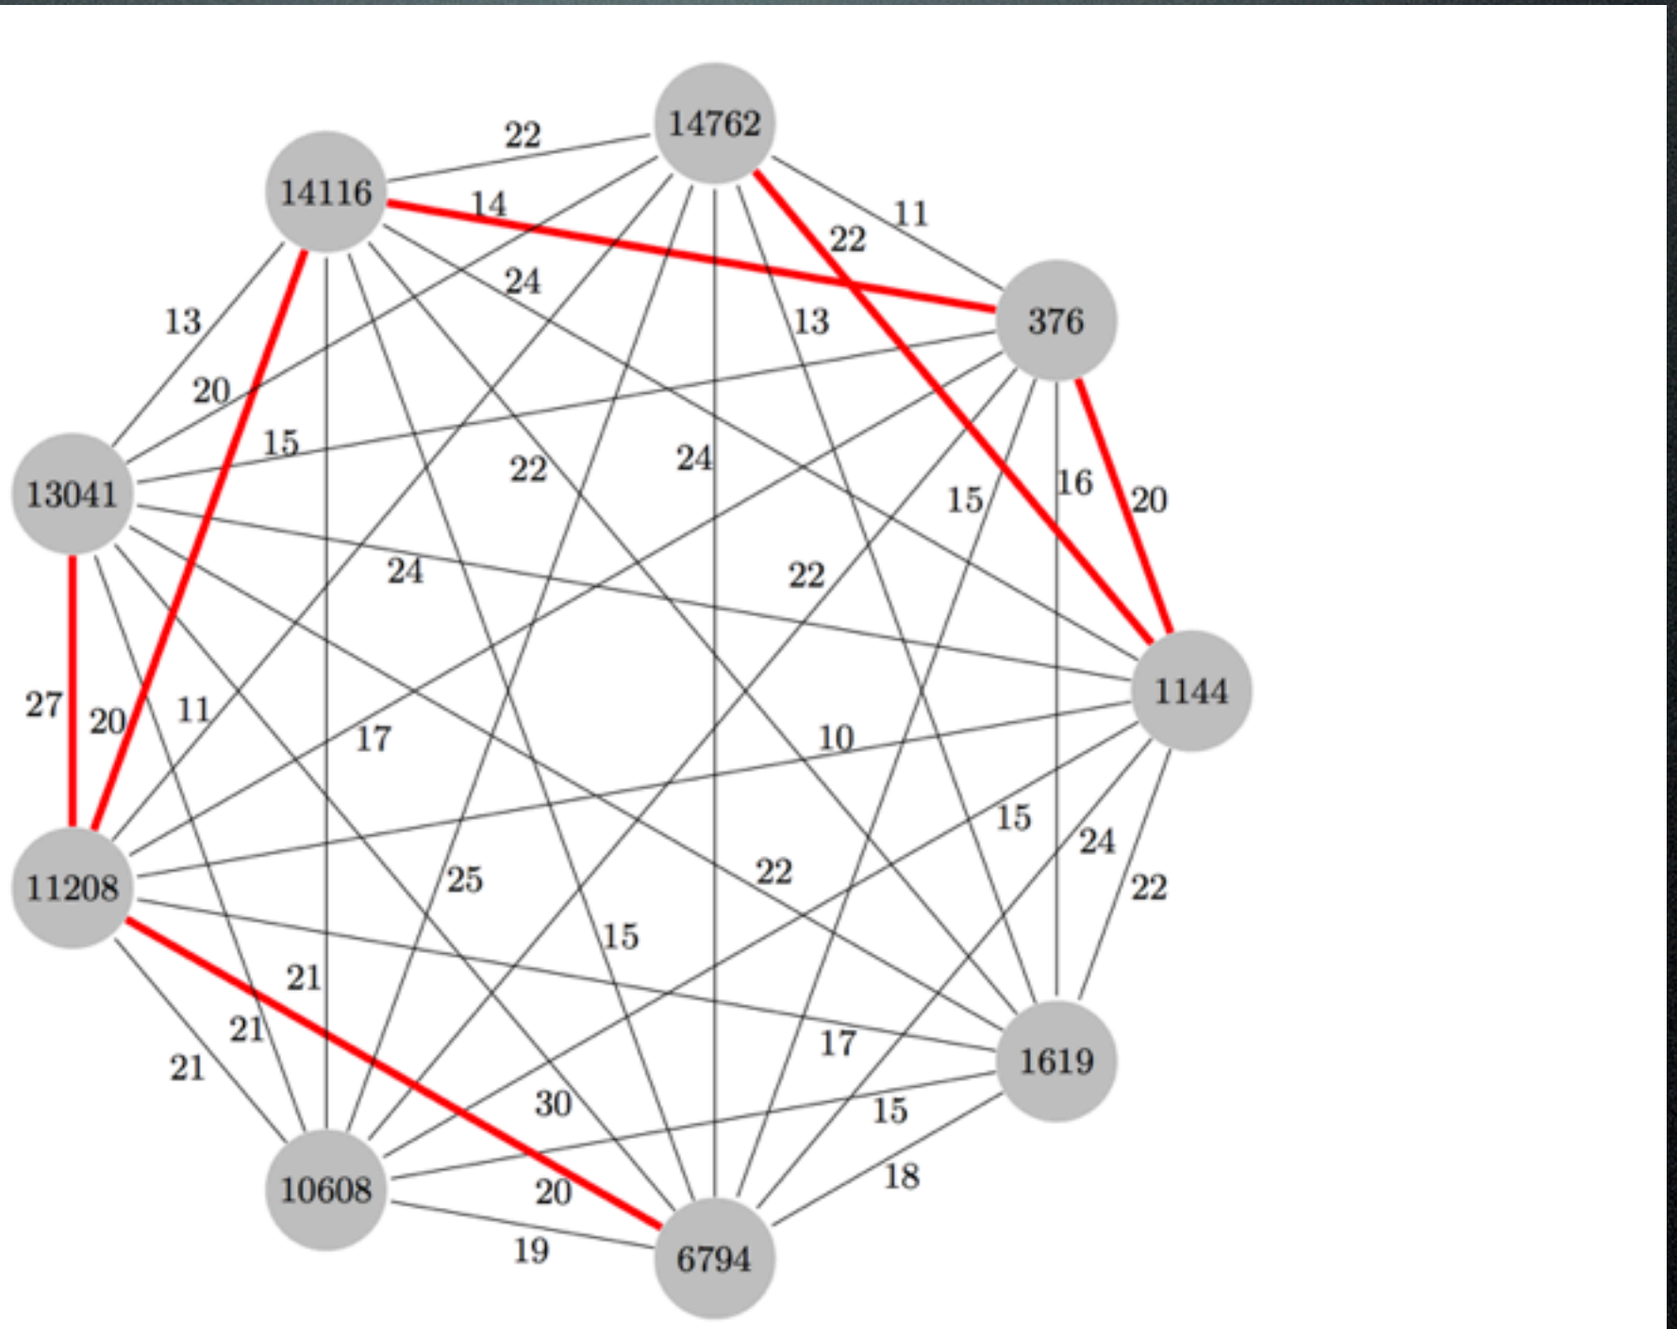

Fig. 6. Connected component consisting of nine vertices, corresponding to primes p_{376}, p_{1144}, ..., p_{14762} (all 512-bit). With labels as in Figure 4, in total 687 X.509 certificates are involved. Six of those certificates have not expired yet, use SHA1 as hash function (as opposed to MD5), and have "`CA=false`"; the red edges correspond to the RSA moduli contained in those six certificates.

Chains

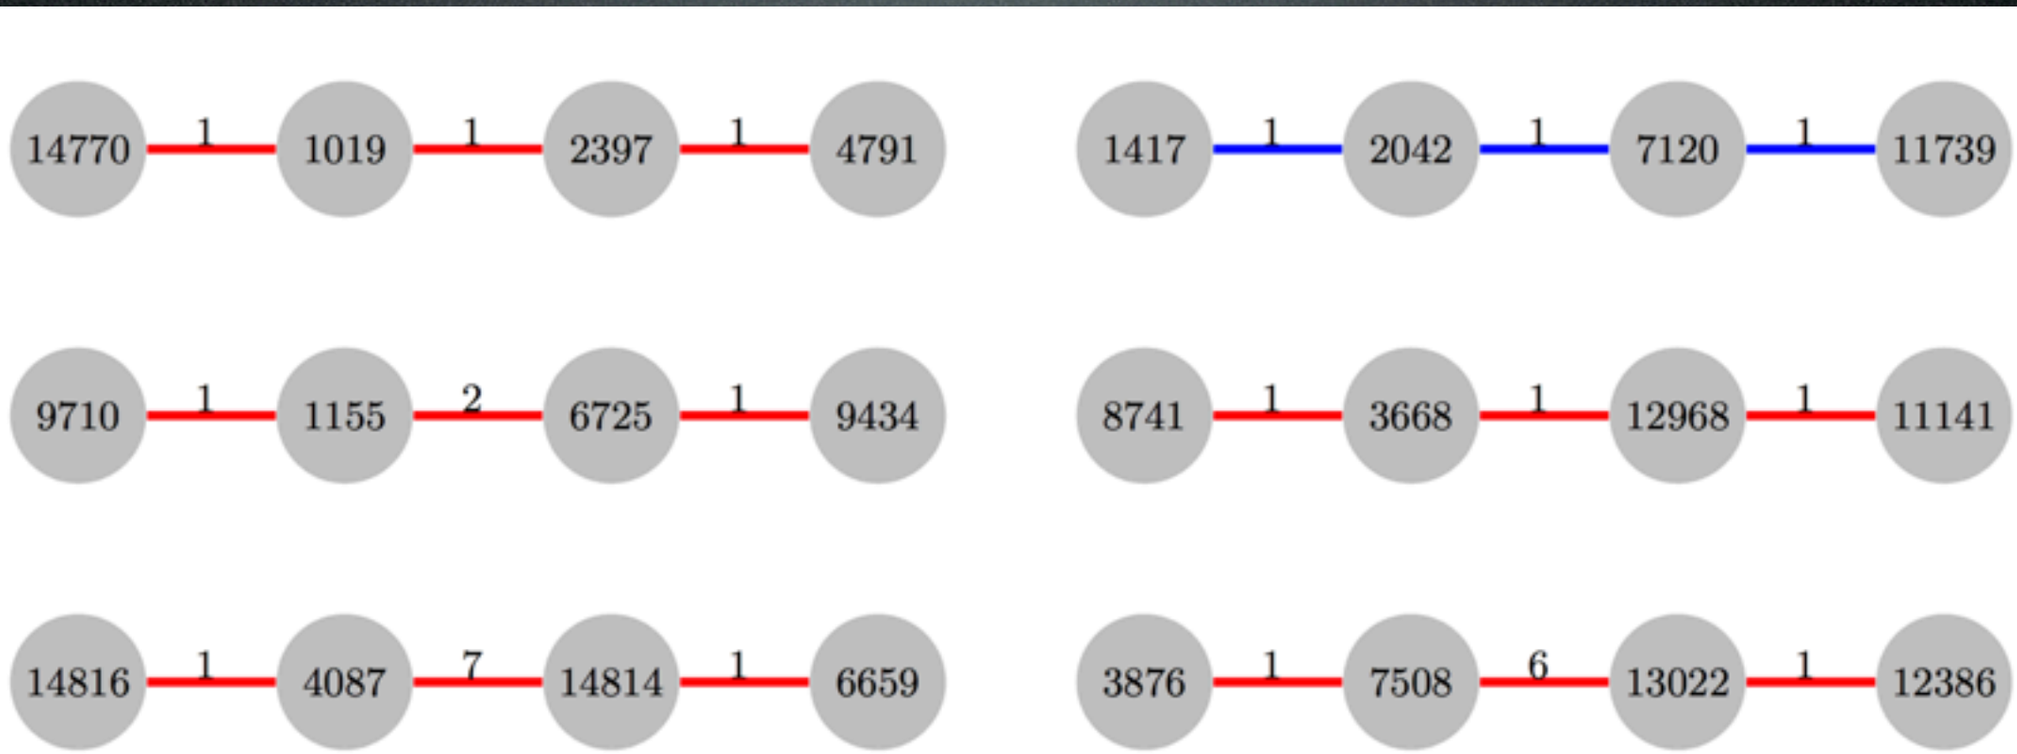

Fig. 5. Six connected components consisting of four vertices, with labels as in Figure 4. The eight primes in the top two components are 512 bits long, the other 16 are 256-bit primes). The red edges correspond to RSA moduli contained in certificates that will not expire anytime soon and that use SHA1 as hash function. The blue ones will expire soon and use MD5.

Discussion

- Bad random number generators will continue to plague the industry.
  - This was not the first instance and won't be the last

Insert Dilbert Commic
http://dilbert.com/strips/comic/2001-10-25/

Discussion

- Bad random number generators will continue to plague the industry.
  - This was not the first instance and ***won't be the last***

*… **generating keys** in the real world for "multiple-secrets" cryptosystems such as RSA is significantly riskier than for "single-secret" ones such as ElGamal or (EC)DSA which are based on Diffie-Hellman.*

- Duplicate keys occur in both
  - Vulnerable to each other
- Only RSA has GCD
  - Complete exposure of private keys

GCD Testing

- Good idea?

  1. Alice creates a key
  2. 10 years pass
  3. Bob creates a key
  4. Testing detects the collision
  5. Alice’s information is compromised

- Alice was an innocent bystander

Discussion: Key Generation

	D-H	RSA
Duplicate Keys	Possible	Possible
Detectable	Compare	Compare
Consequence	Pairwise	Pairwise

Shared Factor		Possible
Detectable		GCD
Consequence		Failure

- Any time there is a detected problem all keys from that particular generator should be revoked.

Conclusion

- Collected 11.7 million public keys
- Recovered thousands of private keys
- Quality RNG are critical
- GCD vulnerability is unique to RSA
  - ECDSA is a very safe alternative

Backup

Key Usage

- DSA has a well known nonce vulnerability
  - Reuse nonce, your keys are divulged
  - Does not affect effect any other keys
    - You can ruin your own day, not someone else’s
- RSA does not require a nonce