

Outline

- 1 Password-based Authenticated Key Exchange
- 2 Game-based Security
- 3 Universal Composability
- 4 Language-based Authenticated Key Exchange

Password-based Authenticated Key Exchange

David Pointcheval

Ecole Normale Supérieure



PKC 2012
Darmstadt, Germany
May 22nd, 2012

Key Exchange Protocols

A fundamental problem in cryptography:

Enable secure communication over insecure channels

A common scenario:

Users encrypt and authenticate their messages using a shared secret key



How to obtain such a shared secret key? → **Key exchange protocols**

Diffie-Hellman Key Exchange

The classical Diffie-Hellman protocol allows such a key exchange: in a finite cyclic group \mathbb{G} , of prime order p , with a generator g

$$\begin{array}{ccc}
 x \xleftarrow{\$} \mathbb{Z}_p, X \leftarrow g^x & \xrightarrow{X} & y \xleftarrow{\$} \mathbb{Z}_p, Y \leftarrow g^y \\
 K \leftarrow Y^x = g^{xy} & \xleftarrow{Y} & K \leftarrow X^y = g^{xy}
 \end{array}$$

No authentication provided

Authenticated Key Exchange

Semantic security / Implicit Authentication:

the session key should be indistinguishable from a random string to all except the expected players

Attacks Examples

On-line Dictionary Attacks The Most Famous Examples

On-line Dictionary Attacks

- The adversary interacts with a player, trying a password
 - In case of success: it has guessed the password
 - In case of failure: it tries again with another password
- If the dictionary has a size N , the adversary wins after $N/2$ attempts

In Practice

- **This attack is unavoidable**
- If the failures for a target user can be detected: the impact can be limited by various techniques (limited number of failures, delays between attempts, ...)
- If the failures cannot be detected (anonymity, no check, ...) the impact can be dramatic

In a finite group \mathbb{G} , of prime order p , with key derivation function \mathcal{K}

EKE: Encrypted Key Exchange [Bellare-Merritt, 1992]

$x \xleftarrow{\$} \mathbb{Z}_p, X \leftarrow g^x$ $X' \leftarrow \mathcal{E}_{pw}(X)$	$\xrightarrow{X \ X'}$ $\xleftarrow{Y \ Y'}$	$y \xleftarrow{\$} \mathbb{Z}_p, Y \leftarrow g^y$ $X \leftarrow \mathcal{D}_{pw}(X')$	DH Key Exchange with flows encrypted under pw
$Y \leftarrow \mathcal{D}_{pw}(Y')$ $k \leftarrow Y^x = g^{xy}$	$\xleftarrow{Y \ Y'}$	$Y' \leftarrow \mathcal{E}_{pw}(Y)$ $k \leftarrow X^y = g^{xy}$	

$K \leftarrow \mathcal{K}(A, B, X, Y, k) \quad K \leftarrow \mathcal{K}(A, B, X', Y', k)$

SPEKE: Simple Password Exponential Key Exchange [Jablon, 1996]

$x \xleftarrow{\$} \mathbb{Z}_p, X \leftarrow g^x$ $k \leftarrow Y^x = g^{xy}$	\xrightarrow{X} \xleftarrow{Y}	$y \xleftarrow{\$} \mathbb{Z}_p, Y \leftarrow g^y$ $k \leftarrow X^y = g^{xy}$	DH Key Exchange with a basis derived from pw
$g \leftarrow \mathcal{G}(A, B, pw)$			

$K \leftarrow \mathcal{K}(A, B, X, Y, k) \quad K \leftarrow \mathcal{K}(A, B, g, X, Y, k)$

PACE: Password Authenticated Connection Establishment

The recent alternative to BAC is PACE: Password Authenticated Connection Establishment
 In the spirit of SPEKE: a generator derived from the password

- With security analyses:
- PACE v1 [Bender-Fischlin-Kuegler, 2009]
 - PACE v2 [Coron-Gouget-Icart-Paillier, 2011]

What does security really mean?

Security Models

- **Game-based Security** [Bellare-P.-Rogaway, 2000]
 - Find-then-Guess
 - Real-or-Random
- Simulation-based Security [Abdalla-Fouque-P., 2005]
- **Universal Composability** [Boyko-MacKenzie-Patel, 2000] [Canetti-Halevi-Katz-Lindell-MacKenzie, 2005]

Where

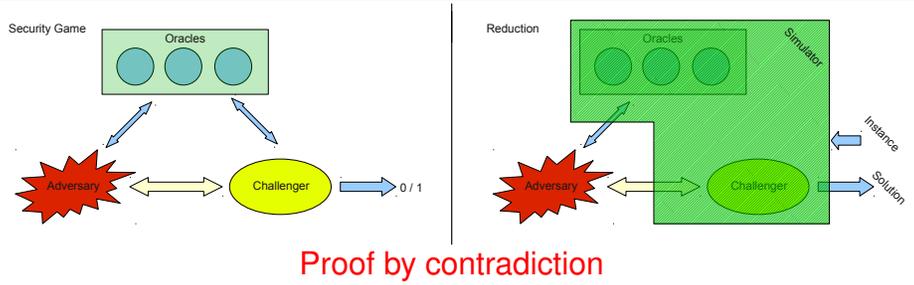
- The adversary controls all the communications: It can create, modify, transfer, alter, delete messages
- Users can participate in concurrent executions of the protocol Instances of the players are denoted A^i and B^i

On-line dictionary attack should be the best attack
 \implies No adversary should win with probability greater than q_S/N
 where $q_S = \#$ Active Sessions and $N = \#$ Dictionary

Game-based Security [Bellare-P.-Rogaway, 2000]

Computational Security Proofs

- a formal security model (security notions)
- a reduction: if one (Adversary) can break the security notions, then one (Simulator + Adversary) can break a hard problem
- acceptable computational assumptions (hard problems)

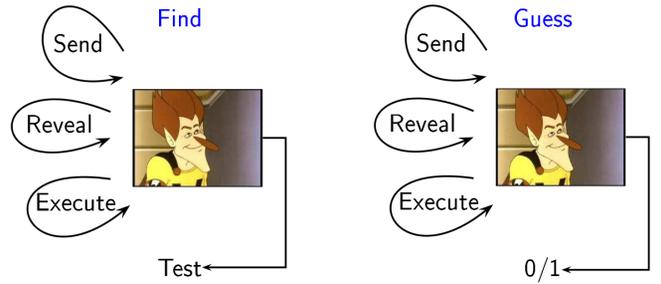


The adversary \mathcal{A} interacts with oracles:

- $Execute(A^i, B^j)$
 \mathcal{A} gets the transcript of an execution between A and B
 It models passive attacks (eavesdropping)
- $Send(U^i, m)$
 \mathcal{A} sends the message m to the instance U^i
 It models active attacks against U^i (active sessions)
- $Reveal(U^i)$
 \mathcal{A} gets the session key established by U^i and its partner
 It models the leakage of the session key, due to a misuse
- $Test(U^i)$ a random bit b is chosen
 - If $b = 0$, \mathcal{A} gets the session key (*i.e.* $Reveal(U^i)$)
 - If $b = 1$, \mathcal{A} gets a random key

Security Game: Find-then-Guess

Secrecy of the key: guess b' of the bit b involved in the Test-query
 Is the obtained key real or random?
Constraint: no Test-query on a trivially known key
i.e. key already revealed through the instance or its partner



$$Adv^{FitG}(\mathcal{A}) = 2 \times \Pr[b' = b] - 1 \leq \frac{O(q_S)}{N} + \text{negl}()$$

Security Games: Advanced Security Notions

- Semantic Security
 The **Find-then-Guess** game models the **secrecy** of the key
 \implies the session key is unknown to the other players
 - What about this secrecy after the corruption of a player?
 - What about the knowledge of the two players?
- Forward Secrecy
 - An additional oracle: $Corrupt(U)$ provides the password pw of the player U to the adversary
 - A new constraint: For any $Test(U^i)$, player U was not corrupted when U^i was involved in its session
- Explicit Authentication
 \implies the session key is **really** known to the two expected players
 The attacker wins the **Explicit Authentication Game** if
 - an instance terminates with a key
 - without exactly one partner having the material to compute the key

Examples **Secure Protocols: EKE-like** Examples **Smooth Projective Hash Functions**

- With both Random Oracles and an Ideal Cipher
- EKE (ROM+ICM) [Bellare–P.–Rogaway, 2000]
 - ⇒ with Forward-Secrecy
 - OEKE (ROM+ICM) [Bresson–Chevassut–P., 2003]
 - ⇒ with Forward-Secrecy and Client-Authentication
 - Formally verified with CryptoVerif [Blanchet, 2012]

- With Random Oracles (and One-time Pad)
- OMDHKE (ROM) [Bresson–Chevassut–P., 2004]
 - ⇒ with Forward-Secrecy and Server-Authentication
 - SPAKE (ROM) [Abdalla–P., 2005]

$$\begin{array}{ccc}
 x \xleftarrow{\$} \mathbb{Z}_p, X \leftarrow g^x & \xrightarrow{X'} & y \xleftarrow{\$} \mathbb{Z}_p, Y \leftarrow g^y \\
 X' \leftarrow X \cdot h^{pw} & & X \leftarrow X' / h^{pw}, k \leftarrow X^y \\
 Y \leftarrow Y' / h^{pw}, k \leftarrow Y^x & \xleftarrow{Y'} & Y' \leftarrow Y \cdot h^{pw} \\
 K \leftarrow \mathcal{K}(A, B, X', Y', pw, k) & &
 \end{array}$$

Quite Simple Scheme

Definition [Cramer–Shoup, 2002] [Gennaro–Lindell, 2003]

Let $\{H\}$ be a family of functions from X to \mathbb{G} and L a subset (language) of this domain X such that, for any point $x \in L$, and a witness w ,

- $H(x) = \text{Hash}_L(hk; x)$, with the secret hashing key hk
- $H(x) = \text{ProjHash}_L(hp; x, w)$, with the public projected key hp

- Hard-Partitioned Subset: L and X hard to distinguish
- Smoothness: if $x \notin L$, $H(x)$ and hp are independent
- Pseudo-Randomness: if $x \in L$, $H(x)$ is pseudo-random, with hp but without a witness w

Examples **Secure Protocols: KOY/GL-like** Advanced Security **Security Game: Real-or-Random** [Abdalla–Fouque–P., 2005]

- With \mathcal{L} = language of the valid commitments of pw
- GL (Standard + CRS) [Gennaro–Lindell, 2003] ⇒ Forward-secrecy

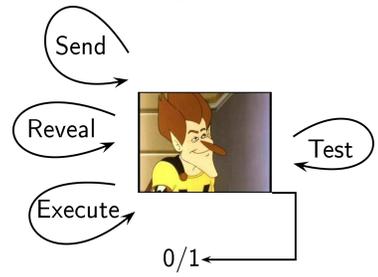
(main steps – more details are required)

$$\begin{array}{ccc}
 C_1 \leftarrow \text{Commit}(pw; r_1) & \xrightarrow{C_1} & C_2 \leftarrow \text{Commit}(pw; r_2) \\
 & \xrightarrow{C_2, hp_1} & hk_1, hp_1 \text{ on } C_1 \\
 hk_2, hp_2 \text{ on } C_2 & \xrightarrow{hp_2} & \\
 \text{ProjHash}(hp_1; C_1, r_1) = H_1 = \text{Hash}(hk_1; C_1) & & \\
 \text{Hash}(hk_2; C_2) = H_2 = \text{ProjHash}(hp_2; C_2, r_2) & & \\
 K \leftarrow H_1 \cdot H_2 & &
 \end{array}$$

- Generalization of the KOY protocol [Katz–Ostrovsky–Yung, 2001]
- With hp_1 and hp_2 independent of C_1 and C_2 resp. ⇒ can be made in **One-Round** only [Katz–Vaikuntanathan, 2011]

Secrecy/independence of all the keys:

- many *Test*-queries on any U^i with the same bit b
- If no key defined by the protocol yet: output \perp
 - If dishonest/corrupted partner: output the real key
 - If player/partner already tested: output the same key
 - If $b = 0$: output the real key
 - If $b = 1$: output a random key



$$Adv^{RoR}(\mathcal{A}) = 2 \times \Pr[b' = b] - 1$$

Security Game: Real-or-Random Game-based Security: Limitations

Semantic Security (Encryption) [Bellare–Desai–Jokipii–Rogaway, 1997]
 Find-then-Guess and Real-or-Random are polynomially equivalent

$$Adv^{RoR}(t, q_T) \leq q_T \times Adv^{FG}(t)$$

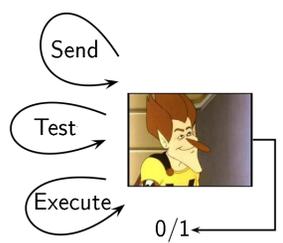
 where q_T is the number of Test-queries

For Password-based Authenticated Key Exchange:

$$Adv^{FG}(t) \leq \frac{\mathcal{O}(q_S)}{N}$$

$$\not\Rightarrow Adv^{RoR}(t, q_T) \leq \frac{\mathcal{O}(q_S)}{N}$$

 \Rightarrow **Much stronger notion**
 No need of Reveal-queries [Abdalla–Fouque–P., 2005]
 \Rightarrow **Much simpler security notion**



- Proven bound: $\mathcal{O}(q_S)/N$, but almost never q_S/N
 \Rightarrow hard to get optimal bound!
 Maybe several passwords can be excluded by each active attack
 - Passwords chosen from pre-determined, known distributions
 - Different passwords are assumed to be independent
 - No security guarantees under arbitrary composition
- \Rightarrow **Universal Composability** more appropriate? [Canetti, 2001]

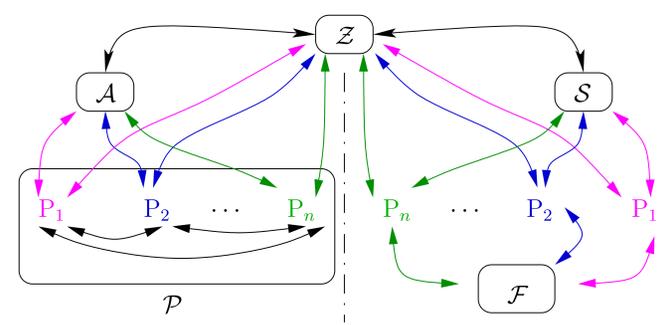
It extends the Simulation-based Security [Boyko–MacKenzie–Patel, 2000]

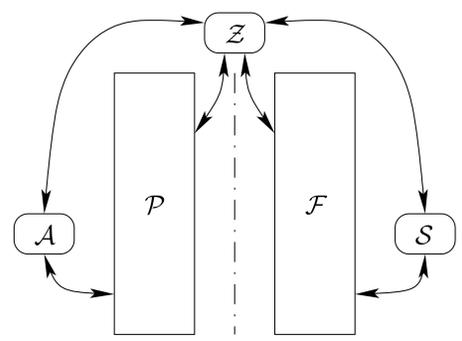
Definition Simulator

Real Protocol
 The real protocol \mathcal{P} is run by players P_1, \dots, P_n ,
 with their own private inputs x_1, \dots, x_n .
 After interactions, they get outputs y_1, \dots, y_n

Ideal Functionality
 An ideal function \mathcal{F} is defined:
 • it takes as input x_1, \dots, x_n ,
 the private information of each player,
 • and outputs y_1, \dots, y_n , given privately to each player
 The players get their results, without interacting:
 this is a “by definition” secure primitive

- \mathcal{P} emulates \mathcal{F} if, for any environment \mathcal{Z} , for any adversary \mathcal{A} ,
 there exists a simulator \mathcal{S} so that, the view of \mathcal{Z} is the same for
- \mathcal{A} attacking the real protocol \mathcal{P}
 - \mathcal{S} attacking the ideal functionality \mathcal{F}





- Everything that the adversary \mathcal{A} can do against \mathcal{P} can be done by the simulator \mathcal{S} against \mathcal{F}
- But the ideal functionality \mathcal{F} is perfectly secure: nothing can be done against \mathcal{F}

Then, nothing can be done against \mathcal{P}

Queries

- `NewSession` = a player joins the system with a password
- `TestPwd` = \mathcal{A} attempts to guess a password (**one** per session)
The adversary learns whether the guess was correct or not
- `NewKey` = \mathcal{A} asks for the session key to be computed and delivered to the player

Corruption-Query

- \mathcal{A} gets the long-term secrets (pw) and the internal state
- \mathcal{A} takes the entire control on the player and plays on its behalf

Corruptions can occur **before the execution**: Static Corruptions
Corruptions can occur **at any moment**: Adaptive Corruptions

PAKE Ideal Functionality

Session Key [Canetti–Halevi–Katz–Lindell–MacKenzie, 2005]

- no corrupted players, same passwords
⇒ same key, randomly chosen
- no corrupted players, different passwords
⇒ independent keys, randomly chosen
- a corrupted player (with the secret from the environment)
⇒ key chosen by the adversary
- correct password guess (`TestPwd`-query)
⇒ key chosen by the adversary
- incorrect password guess (`TestPwd`-query)
⇒ independent keys, randomly chosen

Properties

- The `TestPwd`-query models the on-line dictionary attacks
- The `Corruption`-query includes forward-secrecy

Advantages wrt Game-based Security

- No assumption on the distribution of passwords
- Passwords can be related (it models mistyping)
- Security under arbitrary compositions
⇒ **secure channels**

Game-based Security vs. Universal Compositability Secure Protocols

Game-based Security

In the reduction, the simulator has to emulate the protocol execution **only** up to an evidence the adversary has won ($pw \implies$ not negl.)
 In a global system, the simulation may thus fail as soon as an adversary breaks one of the components whereas other parts could provide protection ($pw \implies$ weak proof!)

UC Security

Handles compositions, but proofs are more complex: the simulator must have an indistinguishable behavior, even when the adversary wins!

In the case of **password-based cryptography**: the adversary can win with non-negligible probability!

In the standard model, with CRS:

- GL^+ (with ZK proofs) [Canetti–Halevi–Katz–Lindell–MacKenzie, 2005]
 \implies Static Corruptions
- With an equivocable/extractable commitment (*bit-by-bit*)
 \implies GL secure against **Adaptive Corruptions** [Abdalla–Chevalier–P., 2009]
- With hp independent of the commitment (*with NIZK*)
 \implies one-round only [Groce–Katz, 2010]
[Katz–Vaikuntanathan, 2011]

With random oracles and an ideal cipher:

OEKE [Abdalla–Catalano–Chevalier–P., 2008]

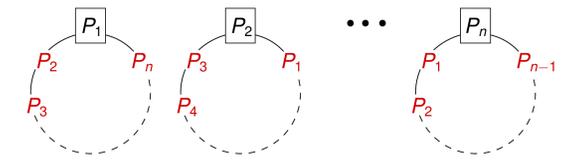
$$\begin{aligned}
 &x \xleftarrow{\$} \mathbb{Z}_p, X \leftarrow g^x \xrightarrow{A, X} y \xleftarrow{\$} \mathbb{Z}_p, Y \leftarrow g^y \\
 &Y \leftarrow \mathcal{D}_{pw}(Y'), K = Y^X \leftarrow \frac{Y'}{Y} \quad Y' \leftarrow \mathcal{E}_{pw}(Y), K = X^Y \\
 &Auth = \mathcal{H}(A, B, X, Y', K) \xrightarrow{Auth} Auth \stackrel{?}{=} \mathcal{H}(A, B, X, Y', K) \\
 &sk = \mathcal{K}(A, B, X, Y', K)
 \end{aligned}$$

\implies First efficient scheme secure against **Adaptive Corruptions**

Weak Authentication: Split Functionality

[Barak–Canetti–Lindell–Pass–Rabin, 2005]

No initial authentication: anybody can join the protocol
 In a multi-party protocol, the adversary can emulate all the other players against one victim, and can do it n times, against the n real players



Split Functionality: initiates a sub-functionality for each sub-session

- Real player P_i : P_i non-corrupted at the beginning
- Adversary on behalf of P_j : P_j corrupted from the beginning

GPAKE: Each sub-session allows to test one password

Limitations of the NewKey-Query

Session Key: NewKey-Query

- a corrupted player \implies key chosen by the adversary
- correct password guess \implies key chosen by the adversary

The NewKey-query is weak

- A lot of control by the adversary: as soon as it controls a player, it controls the key
Key Distribution vs. Key Agreement: Contributiveness
- Not much information leaked to the adversary: whether the protocol succeeds or not
 In practice, the communication continues or stops
 \implies **some information leaks!**

Advanced Security Notions

Contributiveness [Adalla–Catalano–Chevalier–P., 2009]

Initial Definition of the Session Key

- no corrupted players, same passwords \Rightarrow same random key
- corrupted player or correct *TestPwd* \Rightarrow key chosen by \mathcal{A}
- otherwise \Rightarrow independent random keys

With Contributiveness

- at least one non-corrupted player, same passwords \Rightarrow same random key
- all players corrupted \Rightarrow key chosen by \mathcal{A}
- otherwise \Rightarrow independent random keys

It extends to Group protocols, with threshold: (t, n) -Contributiveness
 No player more important than others: \neq key distribution
 Prevents from weak random coins or Trojan horses

Advanced Security Notions

Simpler (but Stronger) Functionality

Queries

- NewSession = a player joins the protocol with a password or \mathcal{A} joins the protocol with a password on behalf of a player $\Rightarrow \mathcal{A}$ impersonates P_i ; it receives the messages for it
- NewKey = \mathcal{A} asks for the session key to be generated
- SendKey = \mathcal{A} asks for the session key to be delivered

NewKey-Query

- the two players are controlled by the adversary \Rightarrow No need to inform anybody: the adversary plays alone!
- Same passwords \Rightarrow same random key – \mathcal{A} informed: OK
- otherwise $\Rightarrow \perp$ – \mathcal{A} informed: NOK

More general \Rightarrow not limited to passwords: Consistent Inputs?

Success Information

The players could learn whether the authentication succeeded

Explicit Authentication

At the Key Delivery time, the player learns: Success or Failure

Together with the Split Functionality:
 the adversary makes a user try a password
 it then learns whether it is correct \Rightarrow similar to *TestPwd*

The adversary should learn this information too (available in practice!)

Successful Agreement

At the Key Computation time, the adversary learns: OK or NOK

In both cases, one can remove the *TestPwd*-query allowing the adversary to join a session with a *NewSession*-query!

LAKE: Ideal Functionality

Queries

- `NewSession` = a player or \mathcal{A} (for a player) joins the protocol with
 - its own language parameters: Pub and $Priv$
 - its partner's language parameters: Pub' and $Priv'$
 - its word w
- `NewKey` = \mathcal{A} asks for the session key to be generated
- `SendKey` = \mathcal{A} asks for the session key to be delivered

Consistent Inputs

The protocol succeeds with the same key if and only if

$$(Pub_a, Priv_a) = (Pub'_b, Priv'_b), \quad (Pub_b, Priv_b) = (Pub'_a, Priv'_a)$$

$$w_a \in L_a(Pub_a, Priv_a), \quad w_b \in L_b(Pub_b, Priv_b)$$

LAKE: Applications

The improved `NewKey`-query
is more powerful/general than the `TestPwd`-query!

LAKE is a quite general framework that includes all the AKE variants:

Particular Instantiations

- $Pub = \emptyset$, $Priv = pw$ and $L(Pub, Priv) = \{Priv\}$
 \implies PAKE (15 group elements exchanged)
 With $Priv = (g^{pw}, h^{pw})$: verifier-based PAKE (29 group elements)
- $Pub = M$, $Priv = vk$, $L(Pub, Priv) = \{\sigma, Verif(Priv, Pub, \sigma) = 1\}$
 \implies Secret Handshake [Balfanz–Durfee–Shankar–Smetters–Staddon–Wong, 2003]
 (43 group elements for Waters Signatures)

Admits efficient instantiations!

LAKE: General Approach

Verification

- $Pub_a = Pub'_b$ & $Pub_b = Pub'_a$: public matching verification
- $Priv_a = Priv'_b$ & $Priv_b = Priv'_a$: implicit matching verification
 \implies as in PAKE
- $w_a \in L_a(Pub_a, Priv_a)$ & $w_b \in L_b(Pub_b, Priv_b)$: implicit verification
 \implies much more complex check!

The GL approach, with advanced Smooth Projective Hash Functions,
allows to implement all these private/implicit checks

(many more details are required)

$$C_1 \leftarrow \text{Commit}(\cdot; r_1) \xrightarrow{C_1} C_2 \leftarrow \text{Commit}(\cdot; r_2)$$

$$\xleftarrow{C_2, hp_1} \quad hk_1, hp_1 \text{ on } C_1$$

$$hk_2, hp_2 \text{ on } C_2 \xleftarrow{hp_2}$$

$$\text{ProjHash}(hp_1; C_1, r_1) = H_1 = \text{Hash}(hk_1; C_1)$$

$$\text{Hash}(hk_2; C_2) = H_2 = \text{ProjHash}(hp_2; C_2, r_2)$$

$$K \leftarrow H_1 \cdot H_2$$

Can be instantiated
under the **DLin** assumption
or the **DDH** assumption

Conclusion

Theoretical Aspects

- Many security models for AKE and PAKE: **Mature Topic**
- Many PAKE candidates:
 - EKE-like protocols are quite efficient, but ideal models
 - GL approach is quite powerful, and reasonably efficient
- LAKE: more general applications, and efficient instantiations

PAKE in Practice

- While appealing, PAKE not really used in practice:
 - IETF RFC 2945 for SRP (no security analysis!)
 - EKE-like: quite efficient but patented \implies not used so far
- EKE Patent expired late 2011 \implies recent IETF RFC 6124

With EKE-like (efficient) or GL-based (fine-grained authentication)
approaches, any situation should find an AKE solution!