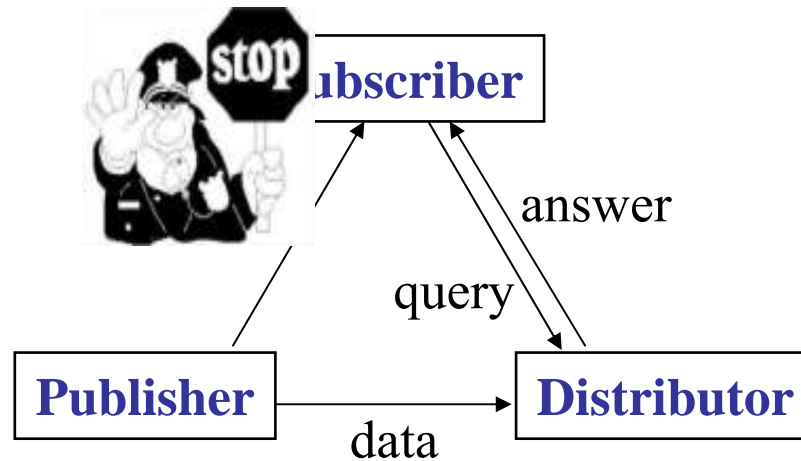


Authenticated Query Processing on Untrusted Servers


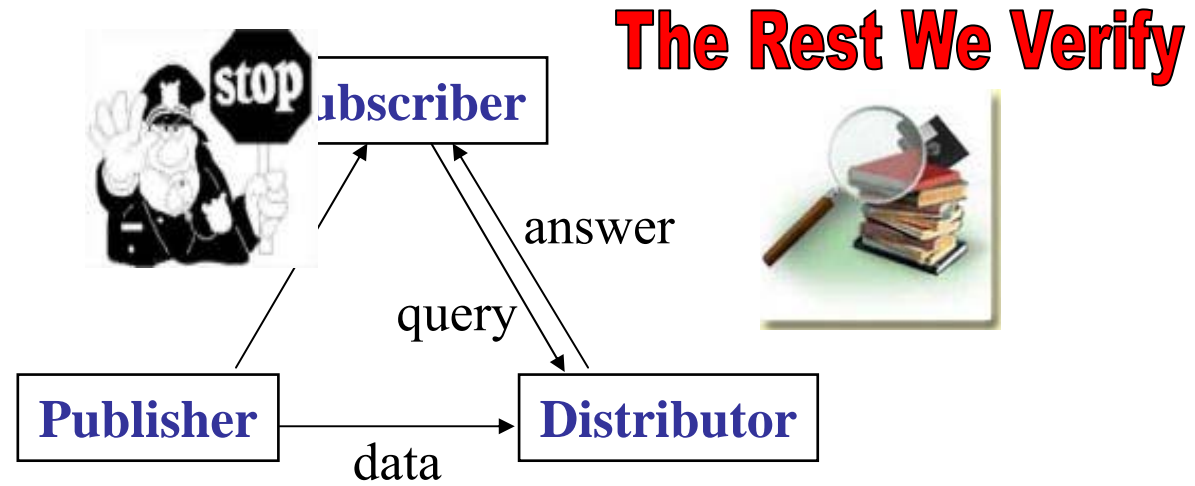
HweeHwa PANG, Ph.D.
School of Information Systems
Singapore Management University

The Case for Untrusted Servers




- Outsourced data management
- Grid, e.g. Condor from University of Wisconsin
- Peer-to-peer platform
- Publish-subscribe / content distribution network

The Challenge



The Challenge



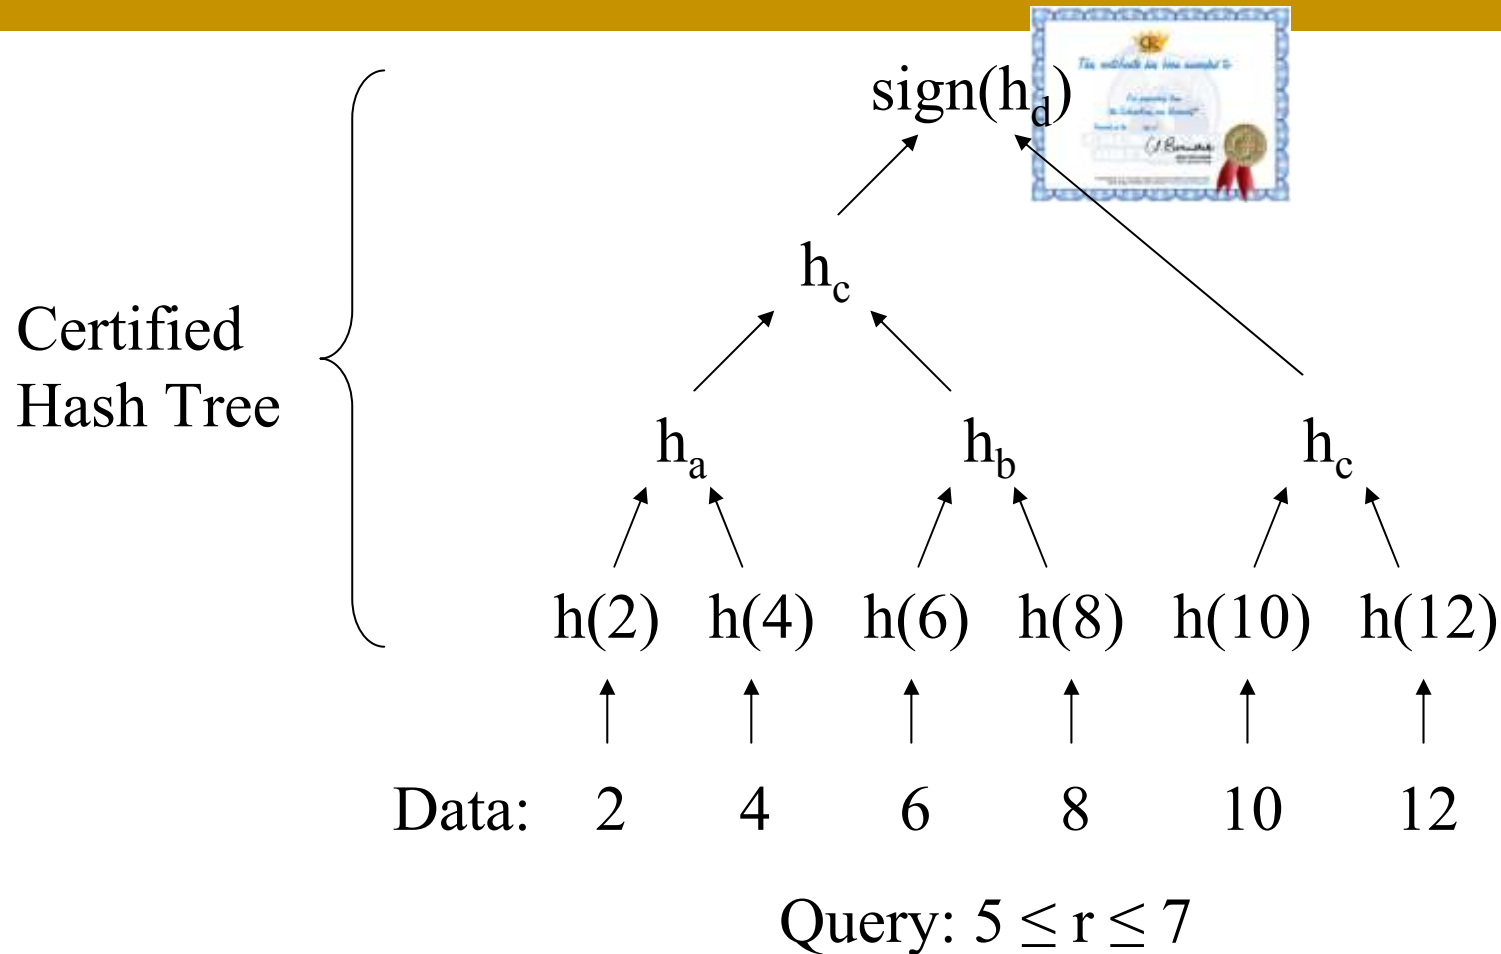
The Truth?
The Whole Truth?
Nothing But The Truth?



Design Objectives:

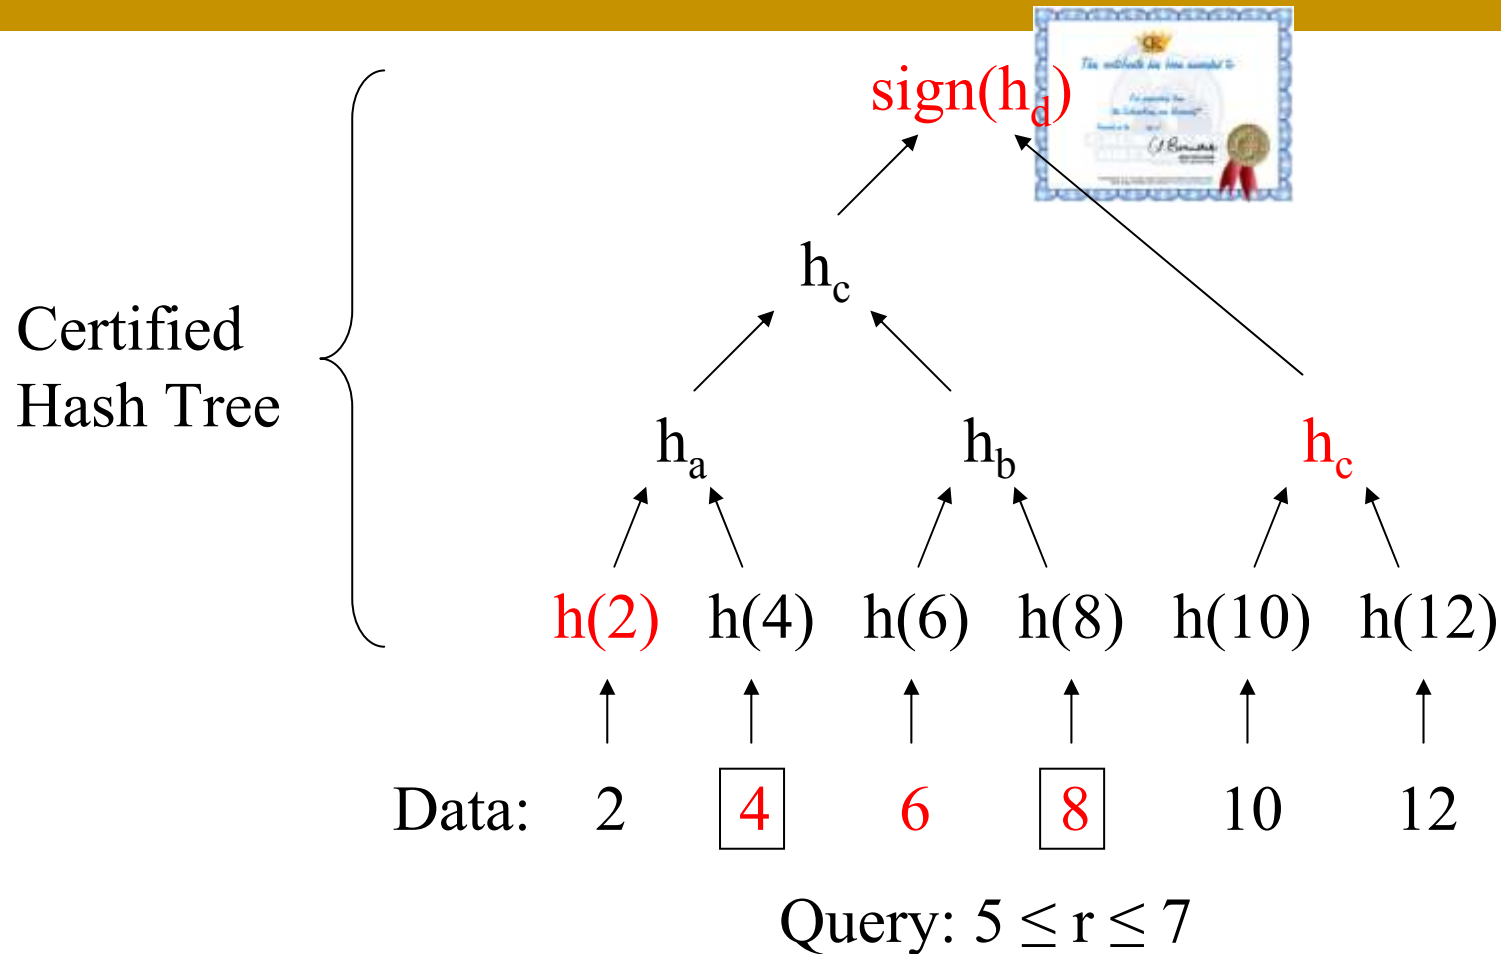
- *Completeness*: No result entry is omitted from the answer
- *Authenticity*: Every entry originated from the publisher
- *Precision*: Minimum information leakage
- *Security*: Computationally infeasible to cheat
- *Efficiency*: Polynomial proof

Completeness is Hard to Prove!



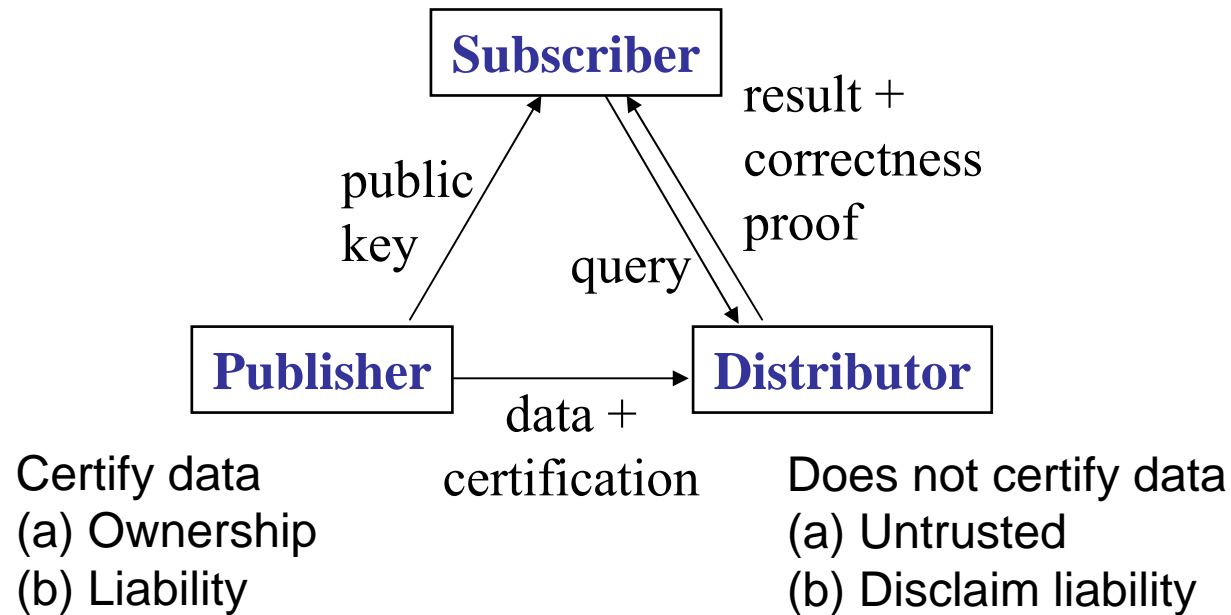
P. Devanbu, M. Gertz, C. Martel, S. Stubblebine, “Authentic Data Publication over the Internet”, 14th IFIP 11.3 Working Conference in Database Security, 2000

Completeness is Hard to Prove!



P. Devanbu, M. Gertz, C. Martel, S. Stubblebine, "Authentic Data Publication over the Internet", 14th IFIP 11.3 Working Conference in Database Security, 2000

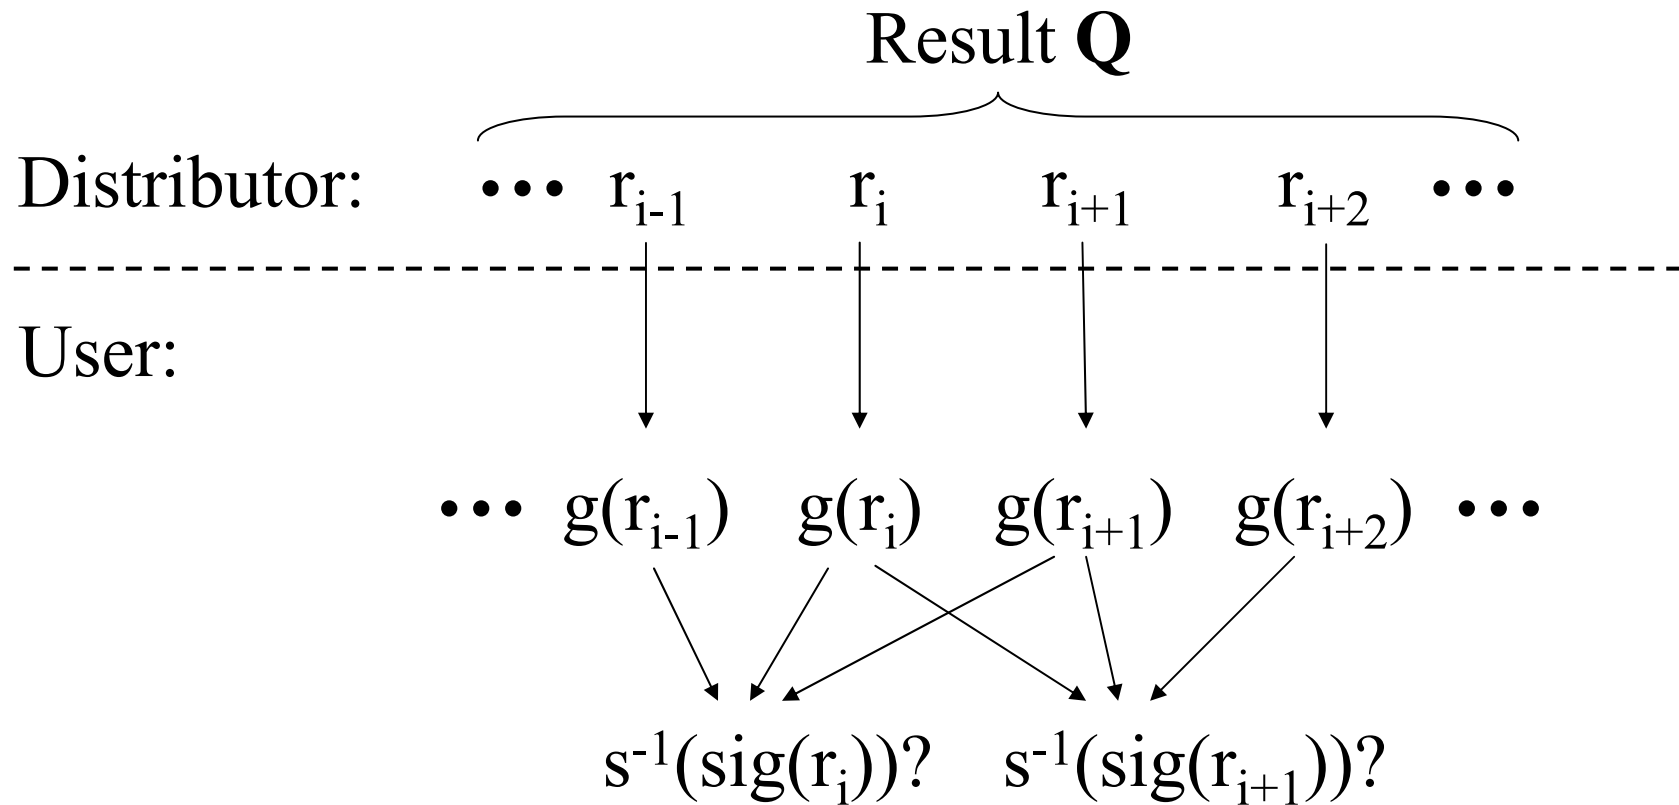
Our Solution



Key Techniques:


- Signature chains: Result entries are contiguous in the data set
- Private boundary proof: Boundary is correct

Signature Chains Ensure Contiguity



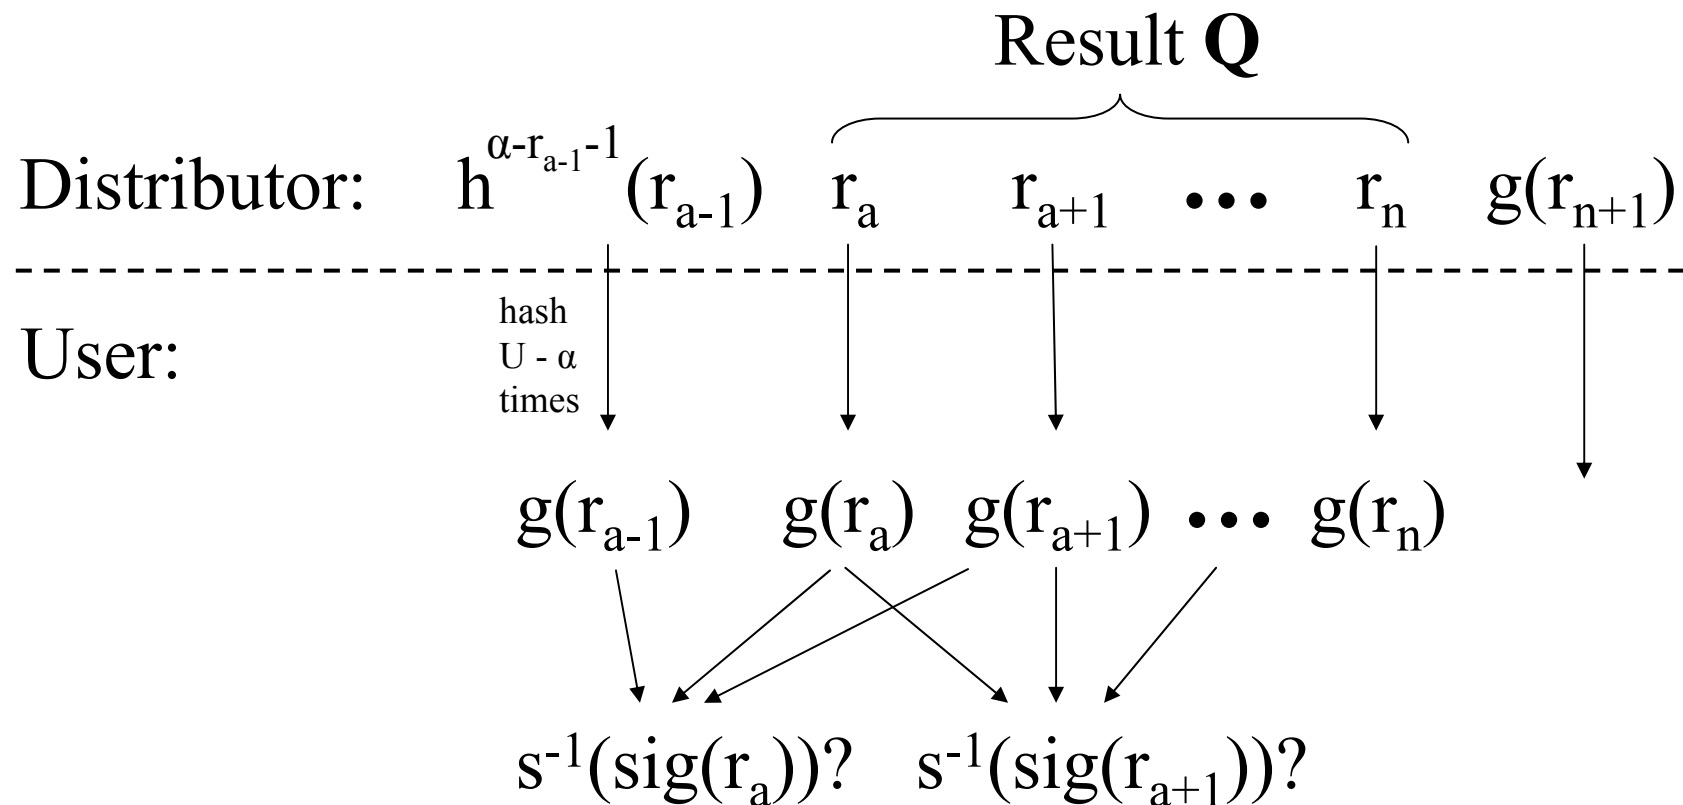
Query: $\alpha \leq r$

Private Boundary Proof



Query: $\alpha \leq r$

Putting the Pieces Together



Query: $\alpha \leq r$



Details can be found in:

H. Pang, A. Jain, K. Ramamritham, K. Tan, “Verifying Completeness of Relational Query Results in Data Publishing”, ACM SIGMOD 2005.