Paper title and authors; where appeared.
Introduction to Modern Cryptography, chapter 6 through section 6.3; by Katz and Lindell.

What is the main problem this paper attacks?
This section examines the feasibility of constructing functions which hide information. Hard-core predicates are examined as a window into some information contained in x which is hidden by a one-way function $f(x)$. A hard-core predicate should be easy to compute given x itself, but infeasible given only $f(x)$.

What solution does the paper propose?
A central result is that given a function f is one-way, $f(x)$ hides the exclusive-or of a random subset of the bits of x. This random exclusive-or process is proposed as a hard-core predicate for any arbitrary one-way function.

What central idea did the authors use to solve it?
The key is to rely on the idea of computational hardness to limit the actions of the adversary. It's easy to invert a one-way permutation given exponential time, by exhausting the domain; however given PPT, we can only invert a one-way function with negligible probability. Similarly, if a PPT adversary can compute a supposed hard-core predicate $gl(x)$ given only $f(x)$, then he can find x itself with non-negligible probability. This would contradict our notion of f as a one-way function.

What is a weakness or limitation of the paper?
It is as of yet unproven whether or not a hard-core predicate exists for any one-way function $f(x)$; the solution is to construct a hard-core predicate for a function g that incorporates both $f(x)$ and enough bits of randomness to obfuscate a significant portion of x.

Why is this paper important?
A major assumption made in the schemes of Chapter 3 was that we were able to construct pseudorandom number generators. Chapter 6 tries to build up to the existence of pseudorandom generators using a lower-level assumption: the existence of one-way functions.