

John the Ripper (JtR)

- ▶ the famous hash cracker
- ▶ primary purpose is to detect weak Unix passwords
- ▶ supports 200+ hash formats (types)
- ▶ supports several kinds of compute devices:
 - ▶ CPU, Xeon Phi
 - ▶ scalar
 - ▶ SIMD: SSE2+/AVX/XOP, AVX2, MIC/AVX-512, AltiVec, NEON
 - ▶ GPU
 - ▶ OpenCL, CUDA
 - ▶ FPGA, Epiphany
 - ▶ currently for bcrypt only

Problems of JtR development

- ▶ scalability of programmers is low due to 200+ formats: sometimes it is hard to apply even 1 optimization to all formats:
 - ▶ important formats get the optimization first
 - ▶ each additional format to optimize eats more time
- ▶ support for each device needs a separate implementation
- ▶ readability degrades when various cases are handled by preprocessor

Aims of john-devkit

- ▶ to separate crypto algorithms, optimizations, and output code for various devices
- ▶ to include optimizations specific for hash cracking and John the Ripper
- ▶ to provide better syntax
- ▶ to retain or improve performance
- ▶ to provide precise control over optimizations
- ▶ to support various devices: CPU, GPU, FPGA
- ▶ to give great output for great input (not for any input)
- ▶ to be simple

Early results

- ▶ john-devkit is not mature
- ▶ 7 formats were implemented separating crypto primitives, optimizations, and device specific code
- ▶ good speeds (over default implementation in JtR):
 - ▶ raw-sha256 +22%
 - ▶ raw-sha224 +20%
 - ▶ raw-sha512 +6%
 - ▶ raw-sha384 +5%
- ▶ bad speeds (but expose interesting features of john-devkit):
 - ▶ raw-sha1 -1%
 - ▶ raw-md4 -11%
 - ▶ raw-md5 -15%
- ▶ optimizations implemented: interleave, vectorization, unroll of loops, early reject, additional batching (loop around algo)
- ▶ all formats got all optimizations without effort

Optimizations

Cracking process

- ▶ we are in attacker's position
- ▶ we have a lot of candidates to try
 - ▶ high parallelism
- ▶ high level algo:
 - ▶ load hashes (once)
 - ▶ generate some candidates
 - ▶ compute hashes (or only parts)
 - ▶ reject most of wrong candidates
 - ▶ check probable passwords precisely (rare case)
 - ▶ generate next batch of candidates and repeat
- ▶ formats are integrated into this process using OOP-like calls over function pointers

Optimizations

- ▶ some optimizations do not affect internals of crypto algorithms in any way and may be added to any algorithm
 - ▶ additional loop around algo to process more candidates in 1 call
 - ▶ OpenMP support
- ▶ other optimizations affect crypto algorithms
 - ▶ vectorization (SIMD)
 - ▶ precomputation
 - ▶ e.g. first few steps in MD*/SHA* for partially changed input
 - ▶ reversal of operations
 - ▶ e.g. last few steps in MD*/SHA*, DES final permutation
 - ▶ loop unrolling
 - ▶ interleaving
 - ▶ bitslicing
 - ▶ and others

Bitslice

- ▶ splits numbers into bits and computes everything through bitwise operations
- ▶ optimization focuses on minimization of Boolean formula (or circuit)
- ▶ Roman Rusakov generated current formulas for S-boxes of DES used in John the Ripper with custom generator
 - ▶ it took 3 months
- ▶ Billy Bob Brumley demonstrated application of simulated annealing algorithm to scheduling of DES S-box instructions
- ▶ so code generation is not new for John the Ripper (not even speaking about C preprocessor)

Other solutions

OpenCL

- ▶ is the first thing I hear when I say about output for both CPU and GPU
- ▶ has quite heavy syntax (based on C)
- ▶ knows nothing about John the Ripper
- ▶ does not have automatic bitslicing

Dynamic formats in John the Ripper

- ▶ were implemented by Jim Fougeron
- ▶ separate crypto primitives from formats
 - ▶ so $\text{md5}(\$p)$ and $\text{md5}(\text{md5}(\$p))$ have one code base
 - ▶ work at runtime
- ▶ john-devkit aims to be able to do similar thing but at compile time and with ability to optimize better
 - ▶ so $\text{md5}(\text{md5}(\$p))$ would get more optimizations (at price of separate code)

C Macros

- ▶ allow to do things, but are not smart
- ▶ an example of loop unroll in Keccak defining all useful variants:

```
[...]
#elif (Unrolling == 3)
#define rounds \
    prepareTheta \
    for(i=0; i<24; i+=3) { \
        thetaRhoPiChiIotaPrepareTheta(i , A, E) \
        thetaRhoPiChiIotaPrepareTheta(i+1, E, A) \
        thetaRhoPiChiIotaPrepareTheta(i+2, A, E) \
        copyStateVariables(A, E) \
    } \
    copyToState(state, A)
#elif (Unrolling == 2)
#define rounds \
    prepareTheta \
    for(i=0; i<24; i+=2) { \
        thetaRhoPiChiIotaPrepareTheta(i , A, E) \
        thetaRhoPiChiIotaPrepareTheta(i+1, E, A) \
    } \
    copyToState(state, A)
```

X-Macro

- ▶ is a tricky way to use macros, most likely with a separate file to be included multiple times:
 - ▶ the file has code with variable parts
 - ▶ these parts are defined before `#include`
- ▶ so `#include` provides a "template engine"
- ▶ example from NetBSD's libcrypt:

```
[...]  
#define HASH_Init SHA1Init  
#define HASH_Update SHA1Update  
#define HASH_Final SHA1Final  
#include "hmac.c"
```


Example of specific instruction

- ▶ separate instruction is used to define state variable, so john-devkit uses a filter to replace initial state with values for SHA-224 having code for SHA-256:

```
def override_state(code, state):
    consts = {}
    for l in code:
        if l[0] == 'new_const':
            consts[l[1]] = l
        if l[0] == 'new_state_var':
            consts[l[2]][2] = str(state.pop(0))
    return code
```


Specific optimization: early reject

- ▶ hashes are long
- ▶ some output values may be computed a bit quicker than others
- ▶ a 32-bit or 64-bit one value is usually enough to reject almost all wrong candidates
- ▶ so john-devkit drops instructions for computation of other output values in main working function and places full implementation into function for precise check of possible password
- ▶ main implementation is vectorized while full implementation is scalar because it checks only 1 candidate

Full Python

- ▶ is available to define algorithms
- ▶ the environment has some objects with overloaded instructions to produce code in IL in a global variable instead of running it right away
- ▶ but any Python code can be used
 - ▶ it is evaluated fully before further steps of code generation
 - ▶ but to produce good output some additional markup may be needed

Thank you!

- ▶ Thank you!
- ▶ code: <https://github.com/AlekseyCherepanov/john-devkit>
- ▶ more technical detail will be on john-dev mailing list
- ▶ my email: lyosha@openwall.com