A compiler approach to Cyber-Security

François de Ferrière
Compilers Expertise Center
STMicroelectronics - Grenoble, France

JAIF 2019, May 23rd, 2019

life.augmented

Securing IOT nodes

- IOT nodes
  - Fast cryptographic primitives for confidentiality, integrity, authenticity & privacy
  - Power and performance constraints
  - Long lifespan
  - Highly connected

- Subject to physical attacks
  - Side Channel Attacks
  - Fault injection attacks
  - Aiming at
    - Obtain sensitive data
    - Bypass protection
    - Reverse engineering

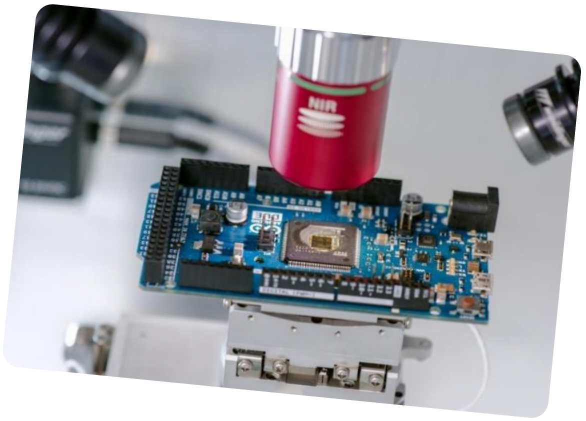

life.augmented

Software-Based Countermeasures

- Source level protections
  - Easy to implement
  - But compiler optimizations tend to remove redundant code
  - Require some implementation tricks and may be difficult to maintain
  - Demoting compiler optimizations results in poor performance and code size

- Assembly level protections
  - The compiler made heavy transformations to reach good performance and code size
  - Difficult to map source code from assembly instructions
  - Difficult to find available resources for adding extra code after aggressive register allocation and code scheduling
  - Higher risk of introducing errors while implementing countermeasures at this level

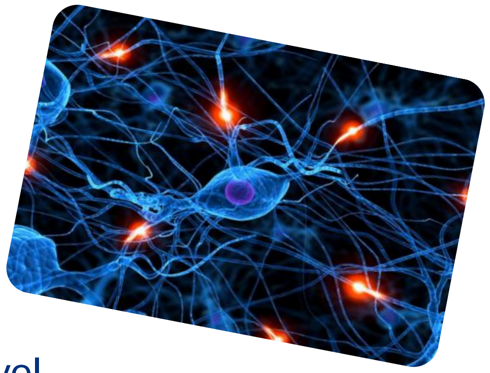

- Aims at protecting against
  - Instruction skip
  - Modification of instructions or data

- A compiler approach
  - Instead of struggling against the compiler, make the compiler work for us
    - No need to modify the source code of an application
    - No need to demote compiler optimizations
  - Security code added by the compiler is part of the code to generate
    - Efficient register allocation and instruction scheduling

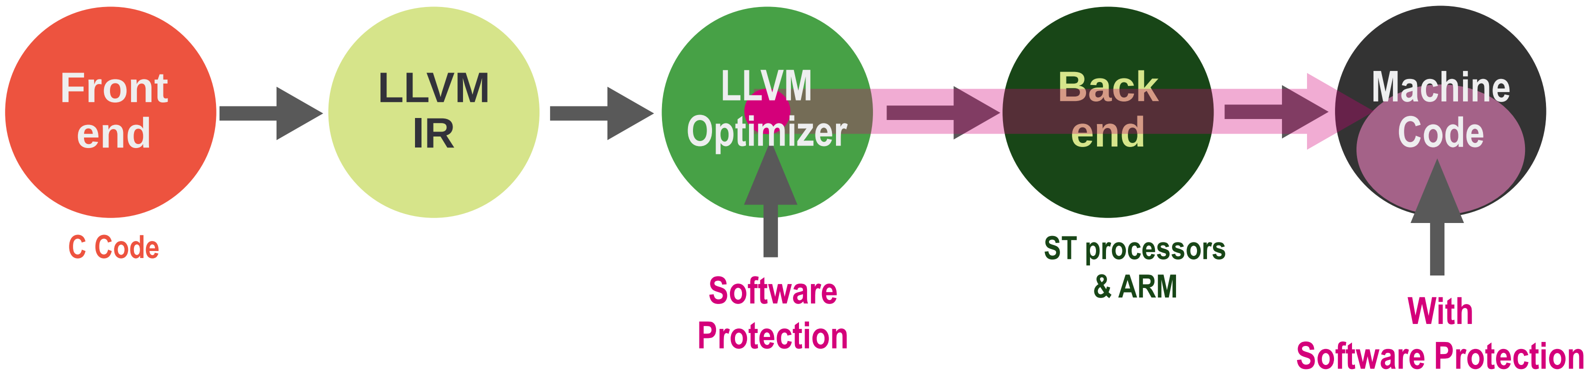

History

- EDDI : Error Detection by Duplicated Instructions in super-scalar processors

  N. Oh, P.P. Shirvani, E.J. McCluskey - IEEE Transactions on Reliability 2002

  - Duplicate instructions and use different registers
  - Duplicate memory locations
  - Check points at side effects

- SWIFT : Software Implemented Fault Tolerance

  G.A. Reis, J. Chang, N. Vachharajani, R. Rangan, D.J. August – CGO 2005

  - Designed to reduce performance and code size impact
  - No duplicated storage, no duplicated loads/stores
  - Control-flow checking

- Fault Model
  - Single fault on any instruction
  - Protection is guaranteed if applied on whole program
  - Memory is protected by hardware (ECC, …)

Introducing LLVM SecSwift

- Our implementation in LLVM: Secure Swift -> SecSwift
  - Abort on fault detection
- SecSwift consists in two different transformations
  - SecSwift Duplicate
    - Duplicate the computation flow inside functions
    - Duplicate parameters and return values on function calls
    - Check the equality of values at synchronization points
  - SecSwift Control-Flow Integrity
    - Branch instructions inside a function
    - Call and return instructions between functions
    - Propagate a signature along control-flow paths
    - Check validity at synchronization points
  - Can be activated independently
  - Combine efficiently and benefit from each other

SecSwift Duplicate

- Duplicate instructions
  - Done on the intermediate representation of the LLVM compiler
  - Check equality at synchronization points (store, return)
  - Counter-measure for instruction skip
- Duplicated instructions go through the backend
  - The compiler will not remove the redundant code
  - The redundant code is fully integrated with the original code for reg-alloc and scheduling

```
int neq = 0, _DUP_neq = 0;
 for (int i = 0, _DUP_i = 0; i < N; i++, _DUP_i++) {
    neq |= input[i] ^ expected[i];
    _DUP_neq |= input[_DUP_i] ^ expected[_DUP_i];
 }
 secswift_trap(i == _DUP_i);
 secswift_trap(neq == _DUP_neq);
```

SecSwift Inter-Procedural DUP

- Parameters and return values duplication on function calls
  - Change calling convention
  - Counter-measure for corruption of parameters and return values
- A new function prefixed with _SECSWIFT_ is created to implement SecSwift IPDUP
  - The original function is kept
  - A dead function elimination pass after SecSwift will remove unused functions

```
<int, int> _SECSWIFT_is_invalid(int *input, int *_DUP_input, size_t N, size_t _DUP_N) {
 ....
 return <neq, _DUP_neq>;
}
```

SecSwift Duplicate

- Duplication is done after optimizations on the LLVM IR
  - Reduces the performance and code size impact of SecSwift
- Use of an intrinsic function to hide copies of variables
  - Generated as an opaque pseudo COPY operation in the Target Machine LLVM IR
  - The register allocator will allocate duplicated variables in different registers
  - Replaced by a real copy instruction after register allocation
- Not all instructions are duplicated
  - Branch instructions are handled by the SecSwift CFG protection
  - Store instructions are not duplicated, since memory is out of the scope of SecSwift
  - Some values are duplicated by a copy of the result of the original instruction
    - On calls and on volatile load instructions
    - On instructions with “undef” operands
- Might have pending caveats
  - Not 100% coverage for now
    - e.g. prologue/epilogue expansion done after LLVM IR

SecSwift Control-Flow

- Control-flow checking: Dynamically checks that branches reach the expected target
  - Counter-measure for fault or skip of branch instructions
  - Based on the property: `A⊕(A⊕B)=B`
  - A static signature is assigned to each basic block: `GSR (General Signature Register)`
  - A dynamic transfer signature is computed on control-flow edges: `RTS (Runtime Transfer Signature)`
  - A check on the signature is inserted at the beginning of basic blocks which have side effect instructions

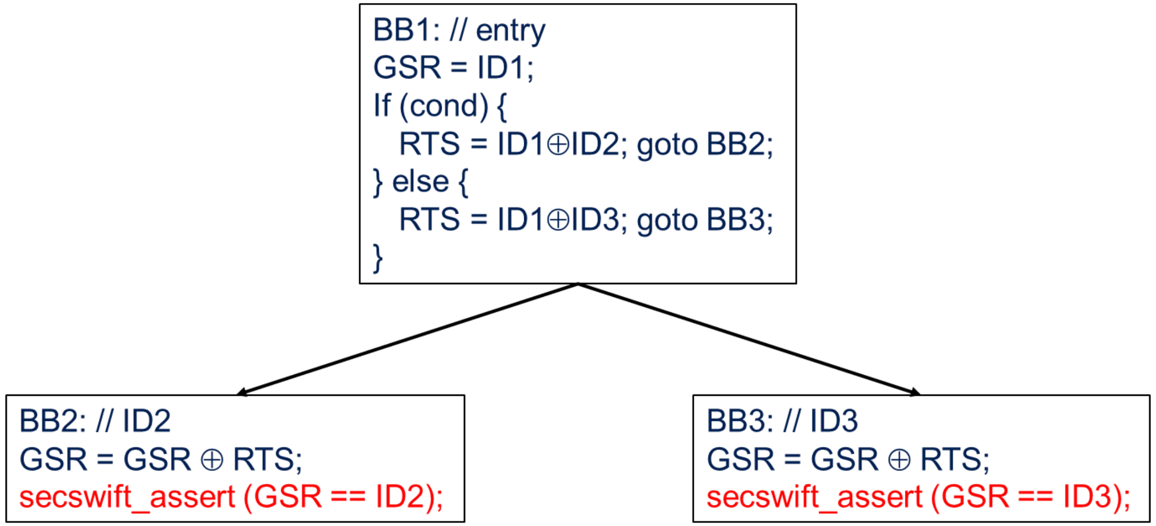

Example 1

```
int GSR = 31155, RTS = 31155 ^ 40106;
for (int i = 0; i < N; i++) {
   GSR ^= RTS;
   neq |= input[i] ^ expected[i];
   RTS = i < N ? 0 : 40106 ^ 642;
}
GSR ^= RTS;
secswift_assert(GSR == 642);
```

Example 2

life.augmented

SecSwift Control-Flow

- Why a XOR ?
  - Mathematical properties
  - Fewer gates, compared to an add or mul
- Why a GSR and RTS ?
  - Creates a chain of updates of the GSR value
  - If one `GSR=GSR⊕RTS` is not executed correctly
    - Because of a fault on the instruction
    - Because of an incorrect control-flow transfer
    - Because of an incorrect value in GSR or RTS
  - The error will be propagated in the next computations of the GSRs
    - No need to insert many checks
      - Only before instructions that do side effects
- GSR serves as a redundant duplicate for the Program Counter

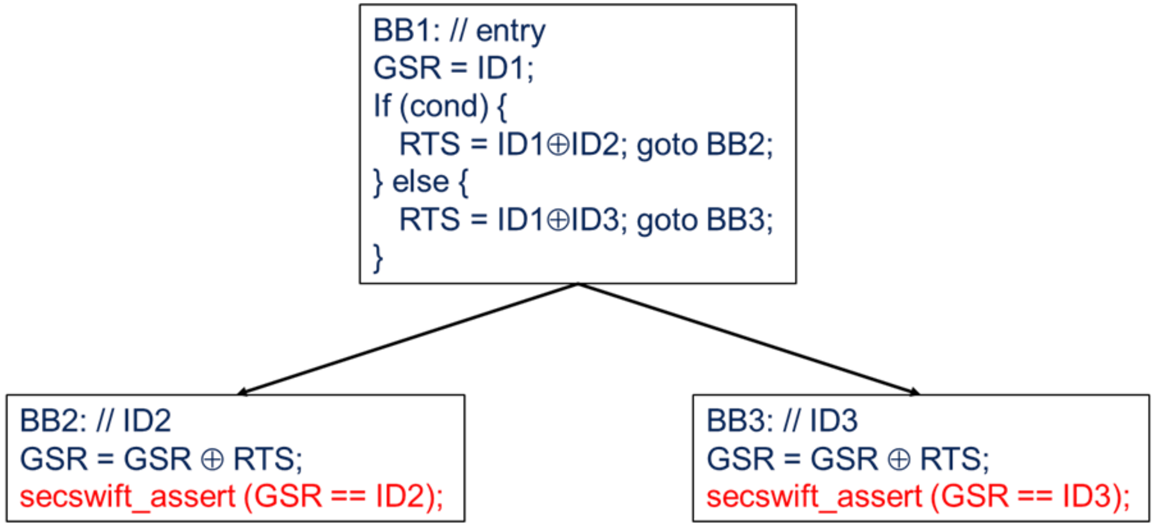

life.augmented

SecSwift Inter-Procedural CFG

- Signatures are statically assigned to functions for which IPCFG has been enabled
  - A hash of the function’s name is used to compute the signatures
  - Two signatures are assigned to each function
    - One for the entry point
    - The other one for all the exit points

- Two parameters, IPGSR and IPRTS, are added on functions protected by IPCFG
  - They replace the GSR and RTS variables on function calls and returns

void f(int *IPGSR, int IPRTS) {
*IPGSR = IDf_e;
…..
g(IPGSR, $IDf_e \oplus IDg_e$);
*IPGSR = *IPGSR $\oplus$ IDg_x;
……
}

void g(int *IPGSR, int IPRTS) {
*IPGSR = *IPGSR $\oplus$ IPRTS;
…….
*IPGSR = *IPGSR $\oplus$ IPRTS $\oplus$ IDg_x;
return;
}

LLVM Implementation Details

- SecSwift passes are implemented at the LLVM IR level
  - Two generic passes
    - One module pass to implement IPDUP and IPCFG transformations
    - One function pass to implement DUP and CFG transformations
  - Added at the very end of the LLVM middle-end passes
  - Do not interfere with general optimizations
  - The pass of Global Dead Function Elimination is run again after SecSwift
    - Eliminate dead functions after the application of SecSwift IPDUP and IPCFG transformations

- Very limited modifications in the target backend
  - We use intrinsic functions and pseudo instructions
    - To prevent copies from being coalesced in the early passes of the Code Generator
    - To generate target dependent code for the SecSwift checks between values
    - They are lowered to real target code before register allocation
  - Support for SecSwift IPDUP on return values
    - Target dependent code on return values duplicated by SecSwift

LLVM Implementation Details

- SecSwift Activation
  - Each SecSwift transformation can be enabled/disabled independently
    - dup : Duplication of the data flow at basic block level
    - cfg : Control-flow integrity checking at basic block level
    - ipdup : Duplication of function parameters and return value
    - ipcfg : Control-flow integrity checking on call and return instructions
- Command line options apply to all functions in a file
  - -fsecswift-…
- Function attributes
  - __attribute__((secswift(…, …)))
  - Override command line options
  - Fine tuning of functions on which SecSwift transformations will be applied

LLVM Implementation Details

- Pragma
  - #pragma secswift(…, …)
  - Override command line options and function attributes
  - Apply to the next single instruction or to the next block of instructions
    - Only ‘dup’ and ‘cfg’ are meaningful
  - Reuse the implementation of the “OpenMP Captured” feature
    - The instructions are outlined into a “captured” function
    - Function attributes are set to the captured function to pass SecSwift options
    - SecSwift is run on captured functions as on other functions
    - The captured function is inlined back into its original function at the end of the SecSwift passes
- SecSwift options are passed from CLANG to LLVM by means of LLVM function attributes
  - Fully validated and functional in LTO mode

Is the generated code more robust ?

- Historically evaluated "by hand"
  - Security experts analyze software protection implemented at source level
  - Then, check in generated code that protections are still there
- The compiler must now be part of the certification process
  - Counter-measures are implemented there
- Tools are needed to improve the evaluation process
  - Simulator with fault injection capability
  - Simple solutions currently in use, based on debugger tools
    - gdb + QEMU on ARM

Is the generated code more robust ?

- Evaluation on a simple string compare function
  - Count the number of successful attacks
    - Success if mcompare returns ‘0’ on different strings
  - Attack is a single skip of an instruction
    - Repeated over every static instruction in the function
    - -O2: 15 instructions, 13% successful attacks
    - -O2 -sec-dup: 53 instructions, 7% succesful attacks
    - -O2 -sec-cfg: 34 instructions, 2% successful attacks
    - --O2 -sec-cfg+dup: 52 instructions, 0% successful attacks
    - -O2 -sec-ipcfg+ipdup: 60 instructions, 0% successful attacks
  - Attack is a clear of a register
    - 100 random pairs instruction/Rx
    - -O2: 15 instructions, 3% successful attacks
    - -O2 -sec-dup: 53 instructions, 3% succesful attacks
    - -O2 -sec-cfg: 34 instructions, 4% successful attacks
    - --O2 -sec-cfg+dup: 52 instructions, 2% successful attacks
    - -O2 -sec-ipcfg+ipdup: 60 instructions, 0% successful attacks

```
int mcompare(unsigned char* s1, unsigned char* s2,
                        unsigned int bytelen) {
  char res = 0;
  int i;
  for (i = 0; i < bytelen; i++) {
    res |= s1[i] ^ s2[i];
  }
  return res;
}
```

How much for this ?

- Evaluation done on ARM Cortex-M0, with options –Oz –flto
  - On a set of 22 benchmarks (eembc, audio/video, dhrystone, coremark, …)
- Performance impact (QEMU instruction count)
  - About 2x slower in average, between 1.5x to 5x
    - Major contribution is -fsecswift-dup
    - -fsecswift-cfg -fsecswift-ipcfg alone is 50% slower in average, 3x at most
    - -fsecswift-ipdup alone has negligible impact
- Code size impact
  - About 3x larger in average, between 1.5x to 4x larger
    - -fsecswift-dup is 2.5x larger in average, 3.5x at most
    - -fsecswift-cfg -fsecswift-ipcfg is 2x larger in average, 3.5x at most
    - -fsecswift-ipdup alone has negligible impact
- Not the whole application code need to be protected
  - Only safety critical application parts
    - Fine scoping through pragmas and function attributes
  - SecSwift impact on performance and code size is comparable to compiling at –O0 without protection

Perspectives

- Continuous race between attacks and countermeasures
  - Fault attacks
    - More and more precise attacks
      - Timing of the attacks
      - Very precise location on a chip
    - Synchronized multiple attacks
  - Countermeasures
    - Protection against skip of multiple instructions has been proposed
    - Add some randomization
      - dead-code
      - random memory location
- No single hardware or software protection, both are needed

Conclusion

- Manually implemented software protection is too limited
  - Sophistication of attacks
  - Complexity of countermeasures
  - Risk on time-to-market
- We provide compilation tools that enable security hardening transformations
  - That would not be reasonably doable by hand – **productivity**
  - That can be local enough to stay limited in resource demand increase - **controllability**
  - That can be global enough to treat arbitrary code bases - **scalability**
  - That play well together - **composability**
  - That are semantically correct for already semantically correct code – **soundness**
- New roles for the security experts
  - Propose new or adapted software counter-measures
  - Validate the counter-measures in the compiler rather than in the final application code
  - Determine which counter-measure are needed on which part of an application

Thanks for your attention

francois.de-ferriere@st.com

life.augmented