

# Lightweight, Multi-Stage, Compiler-Assisted Application Specialization (LMCAS)

Mohannad Alhanahnah and Somesh Jha  
University of Wisconsin-Madison





# Demands for Tiny & Specialized Utilities



**BusyBox**





# Tiny Utilities Generated Manually



# BusyBox

# Our Goal



# BusyBox

# Compilers are Powerful



01

- fauto-inc-dec
- fbranch-count-reg
- fcombine-stack-adjustments
- fcompare-elim
- fcprop-registers
- fdce
- fdefer-pop
- fdelayed-branch
- fdse
- fforward-propagate
- fguess-branch-probability
- fif-conversion
- fif-conversion2
- finline-functions-called-once
- fipa-modref
- fipa-profile
- fipa-pure-const
- fipa-reference
- fipa-reference-addressable
- fmerge-constants
- fmove-loop-invariants
- fmove-loop-stores
- fomit-frame-pointer
- freorder-blocks
- fshrink-wrap
- fshrink-wrap-separate
- fsplit-wide-types
- fssa-backprop
- fssa-phiopt
- ftree-bit-ccp
- ftree-ccp
- ftree-ch
- ftree-coalesce-vars
- ftree-copy-prop
- ftree-dce
- ftree-dominator-opts
- ftree-dse
- ftree-forwprop
- ftree-fre
- ftree-phi-prop
- ftree-pta
- ftree-scev-cprop
- ftree-sink
- ftree-slsr
- ftree-sra
- ftree-ter
- funit-at-a-time

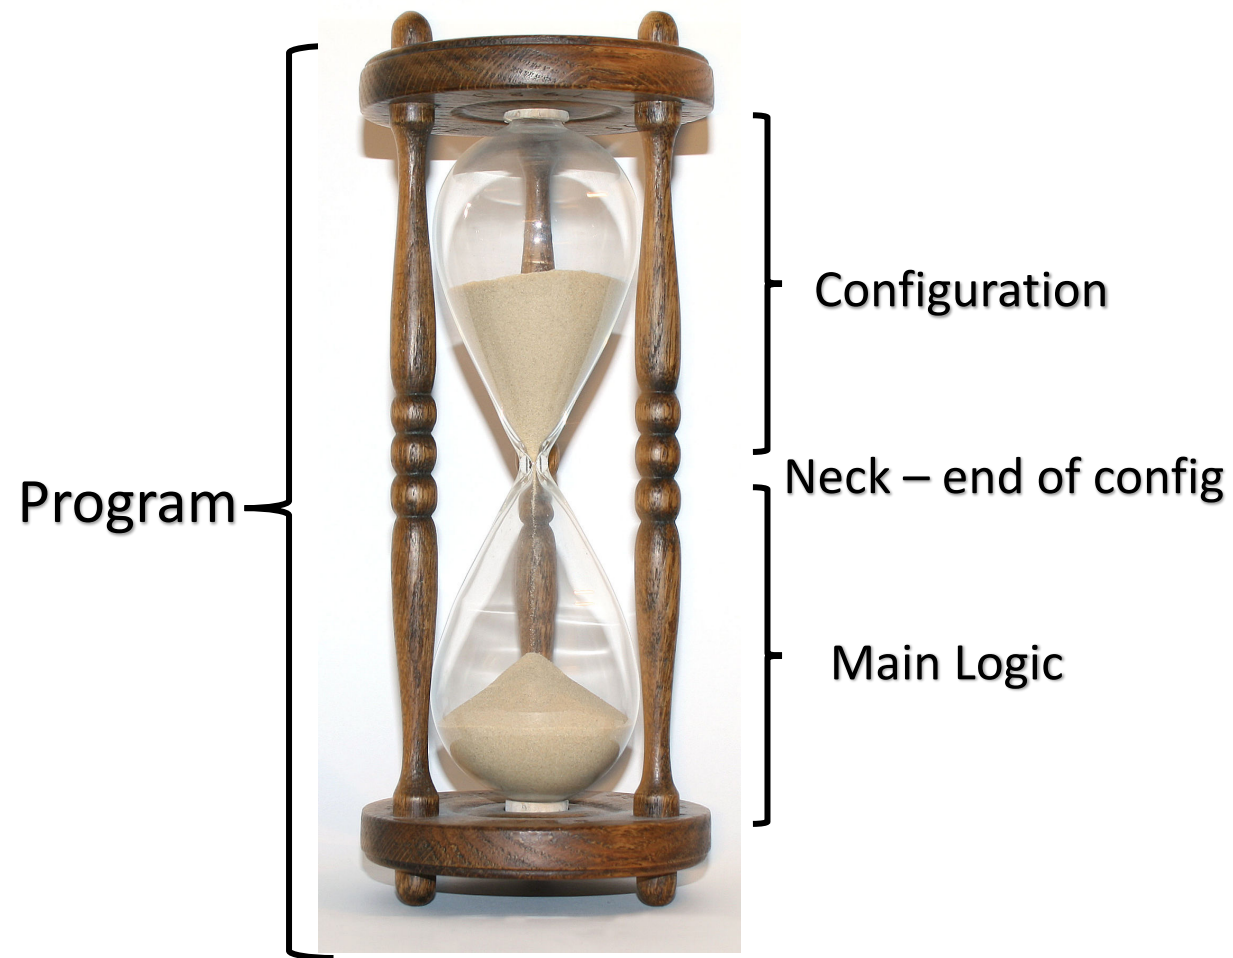
02

- falign-functions
- falign-jumps
- falign-labels
- falign-loops
- fcaller-saves
- fcode-hoisting
- fcrossjumping
- fcse-follow-jumps
- fcse-skip-blocks
- fdelete-null-pointer-checks
- fdevirtualize
- fdevirtualize-speculatively
- fexpensive-optimizations
- ffinite-loops
- fgcse
- fgcse-lm
- fhoist-adjacent-loads
- finline-functions
- finline-small-functions
- findirect-inlining
- fipa-bit-cp
- fipa-cp
- fipa-icf
- fipa-ra
- fipa-sra
- fipa-vrp
- fisolte-erroneous-paths-dereference
- fira-remat
- foptimize-sibling-calls
- foptimize-strlen
- fpartial-inlining
- fpeephole2
- freorder-blocks-algorithm=stc
- freorder-blocks-and-partition
- freorder-functions
- frerun-cse-after-loop
- fschedule-insns
- fschedule-insns2
- fsched-interblock
- fsched-spec
- fstore-merging
- fstrict-aliasing
- fthread-jumps
- ftree-builtin-call-dce
- ftree-pre
- ftree-switch-conversion
- ftree-tail-merge
- ftree-vrp

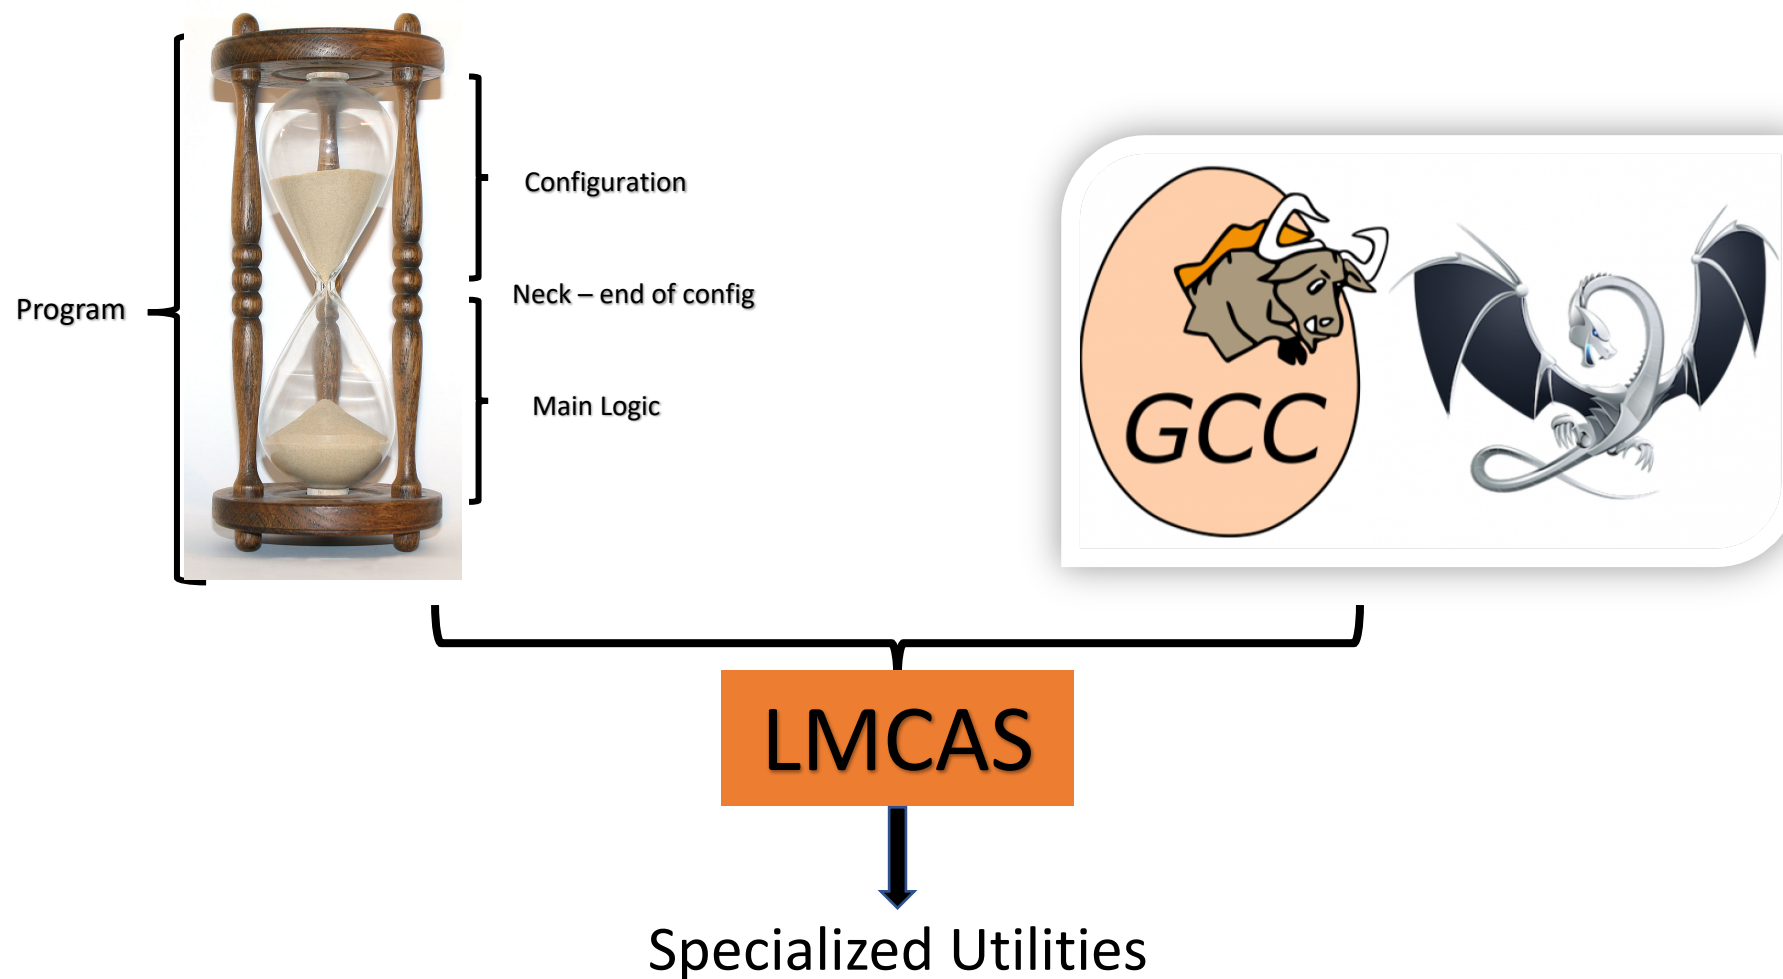
03

- fgcse-after-reload
- fipa-cp-clone
- floop-interchange
- floop-unroll-and-jam
- fpeel-loops
- fpredictive-commoning
- fsplit-loops
- fsplit-paths
- ftree-loop-distribution
- ftree-loop-vectorize
- ftree-partial-pre
- ftree-slp-vectorize
- funswitch-loops
- fvect-cost-model
- fvect-cost-model=dynamic
- fversion-loops-for-strides

# Disciplined Software Development



# LMCAS Pillars



# Agenda

- Use Case
- Introducing LMCAS
- Questions
- Demo




# Survey

- Do you follow Disciplined Software Development?

**YES**

**NO**

Program



Configuration

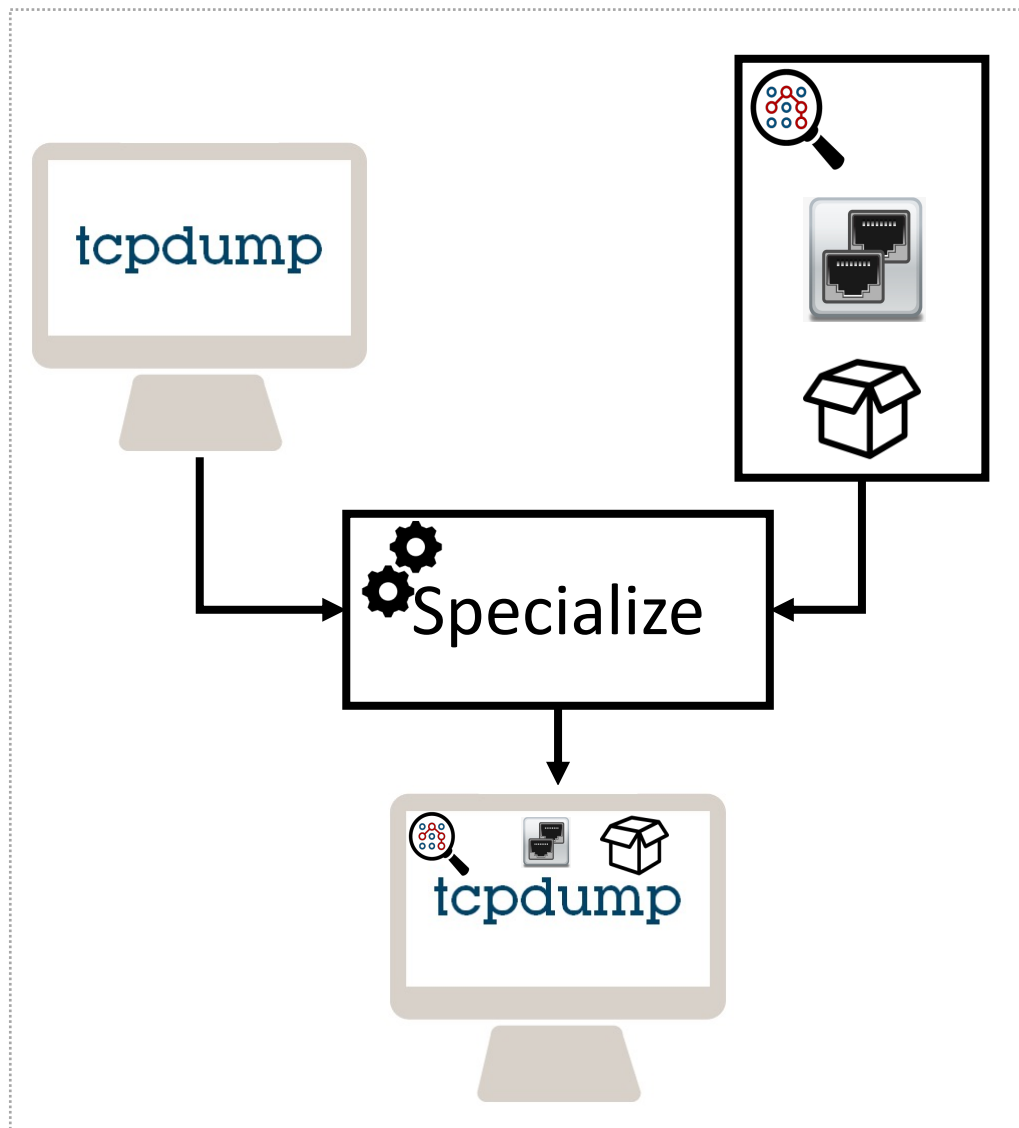
Neck – end of config

Main Logic


# Agenda

- Use Case
- Introducing LMCAS
- Questions
- Demo

# Use Case (Network Monitoring)



# Use Case (Network Monitoring)



Specialization = High Efficiency



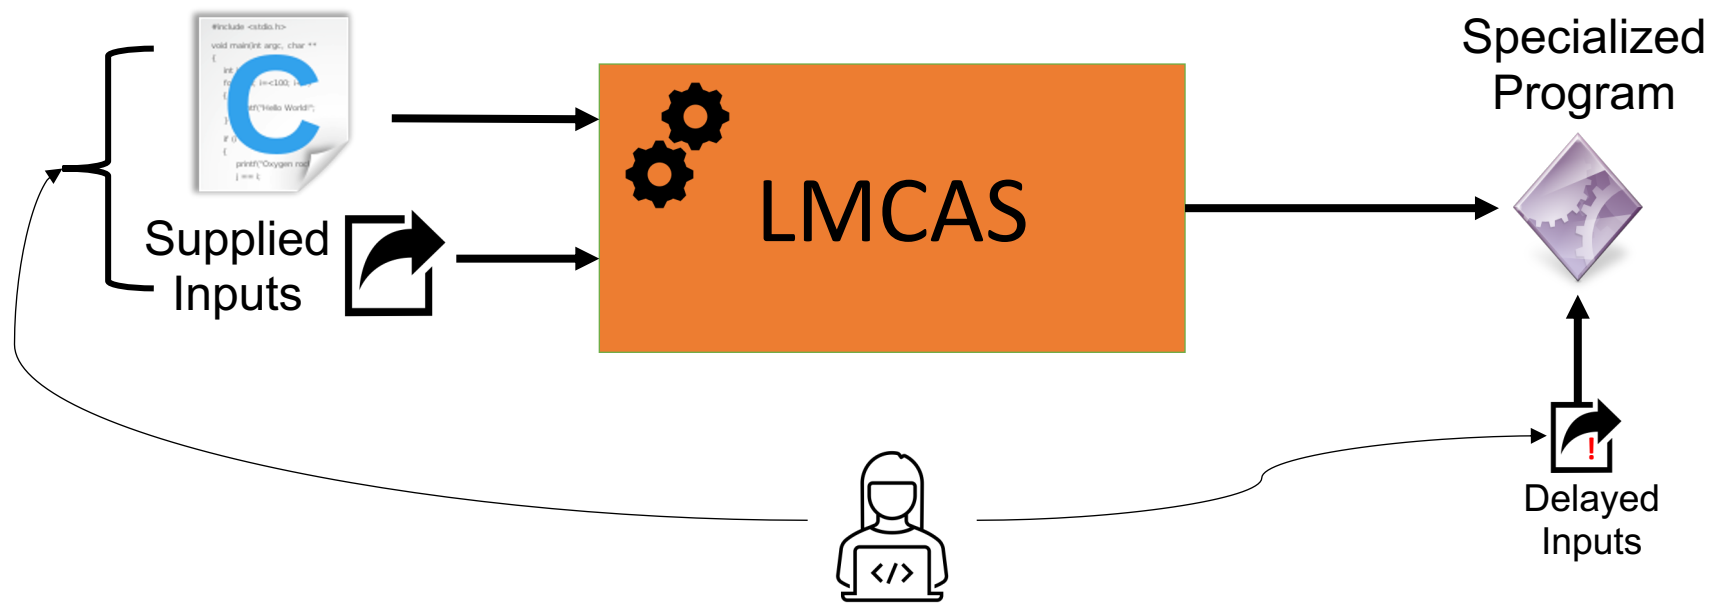
# Other Use Cases



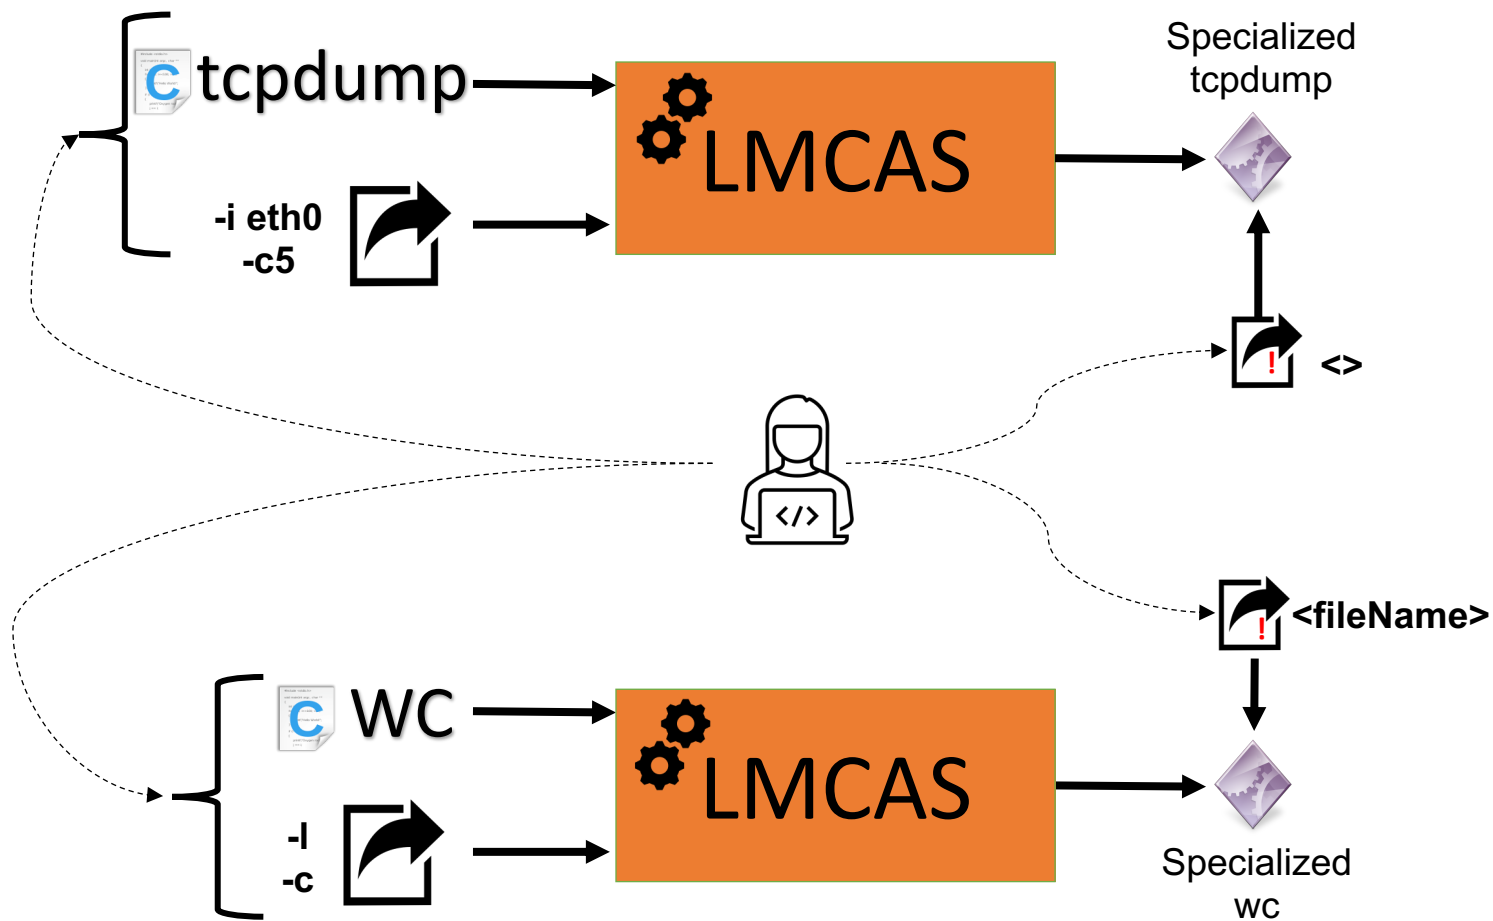
# Agenda

- Use Case
- **Introducing LMCAS**
- Questions
- Demo

# LMCAS Workflow (Partial Evaluation)



# LMCAS Workflow - Examples





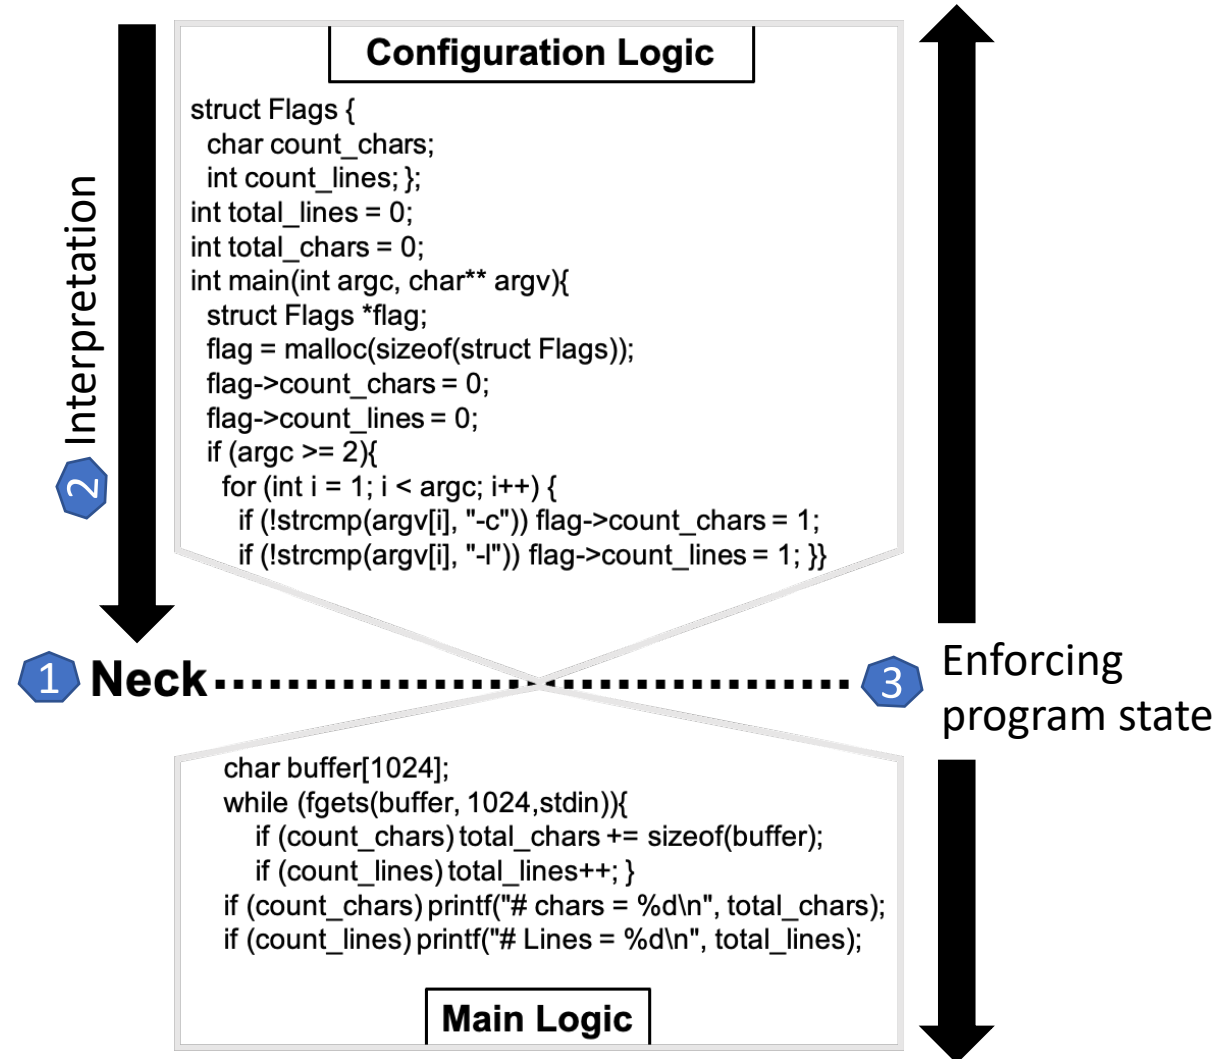
# Illustrative Example

- scaled-down version of the *wc* utility
  1. Line count
  2. Char count

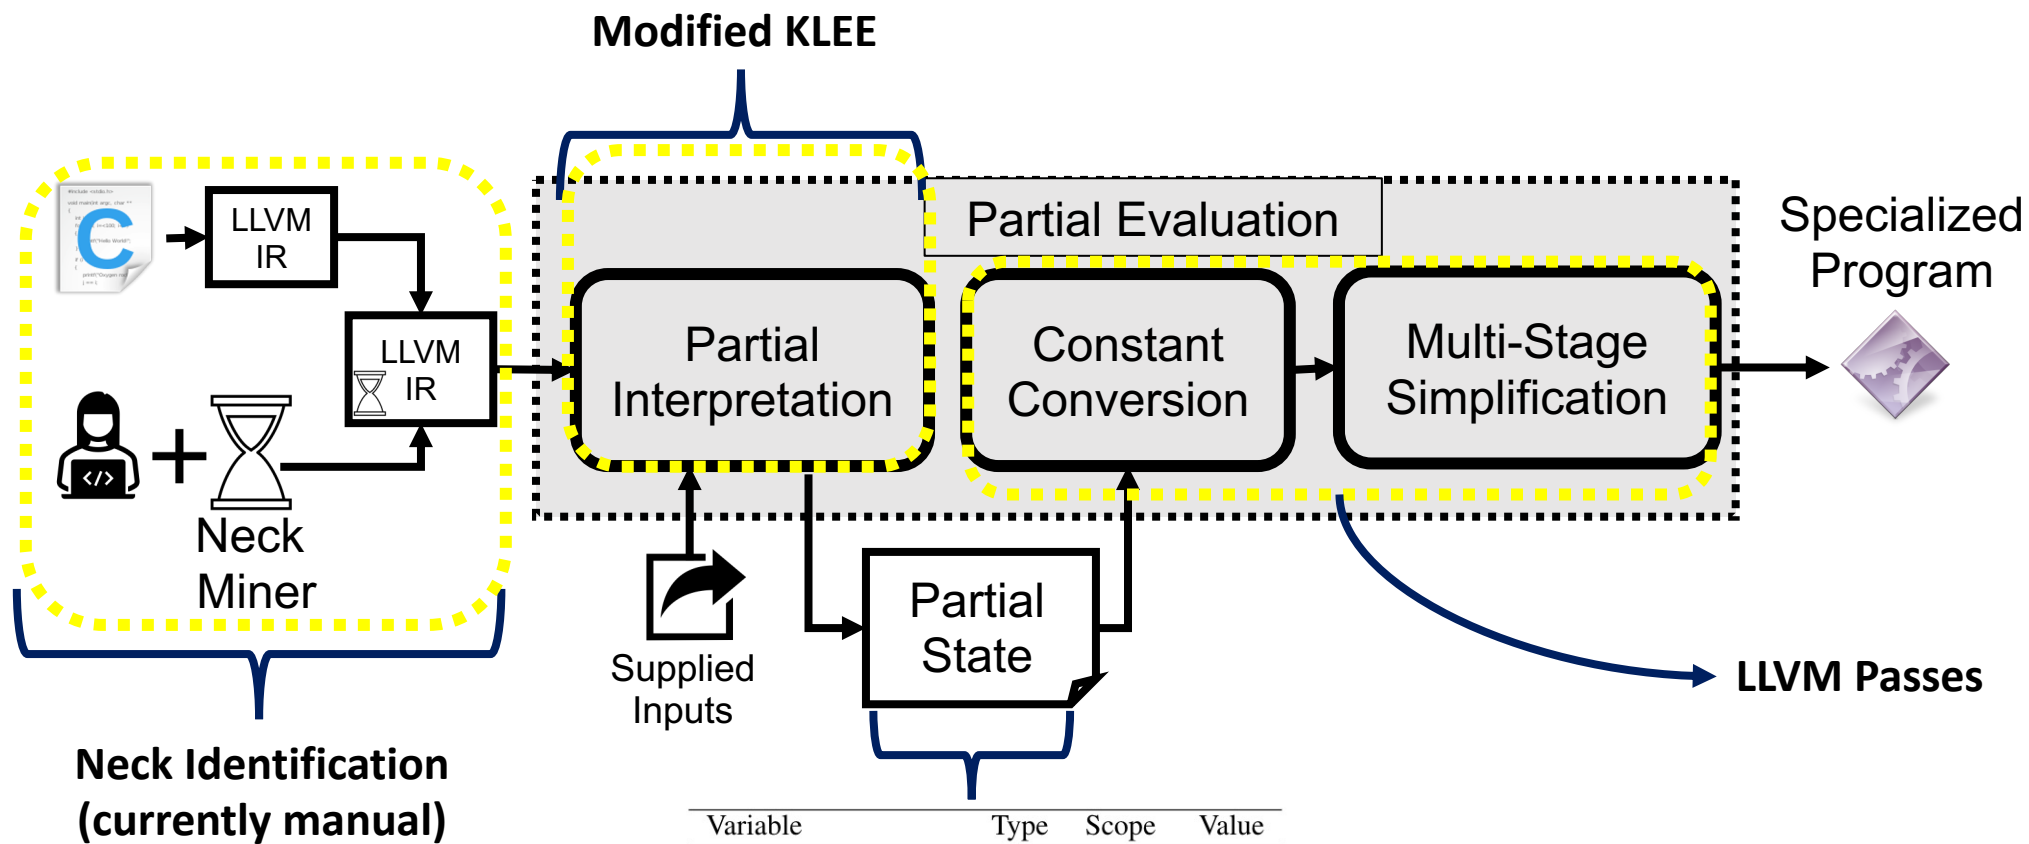
```
1 struct Flags {
2     char count_chars;
3     int count_lines; };
4 int total_lines = 0;
5 int total_chars = 0;
6 int main(int argc, char** argv){
7     struct Flags *flag;
8     flag = malloc(sizeof(struct Flags));
9     flag->count_chars = 0;
10    flag->count_lines = 0;
11    if (argc >= 2){
12        for (int i = 1; i < argc; i++) {
13            if (!strcmp(argv[i], "-c")) flag->count_chars = 1;
14            if (!strcmp(argv[i], "-l")) flag->count_lines = 1;
15        }
16    }
17    char buffer[1024];
18    while (fgets(buffer, 1024, stdin)){
19        if (flag->count_chars) total_chars += decodeChar(buffer);
20        if (flag->count_lines) total_lines++;}
21    if (flag->count_chars) printf("#Chars= %d", total_chars);
22    if (flag->count_lines) printf("#Lines= %d", total_lines);
23 }
```

# LMCAS Approach

```
1 struct Flags {
2   char count_chars;
3   int count_lines; };
4 int total_lines = 0;
5 int total_chars = 0;
6 int main(int argc, char** argv){
7   struct Flags *flag;
8   flag = malloc(sizeof(struct Flags));
9   flag->count_chars = 0;
10  flag->count_lines = 0;
11  if (argc >= 2){
12    for (int i = 1; i < argc; i++) {
13      if (!strcmp(argv[i], "-c")) flag->count_chars = 1;
14      if (!strcmp(argv[i], "-l")) flag->count_lines = 1;
15    }
16  }
17  char buffer[1024];
18  while (fgets(buffer, 1024, stdin)){
19    if (flag->count_chars) total_chars += decodeChar(buffer);
20    if (flag->count_lines) total_lines++;}
21  if (flag->count_chars) printf("#Chars= %d", total_chars);
22  if (flag->count_lines) printf("#Lines= %d", total_lines);
23 }
```



# LMCAS Architecture



| Variable          | Type | Scope  | Value |
|-------------------|------|--------|-------|
| total_lines       | int  | Global | 0     |
| total_chars       |      |        | 0     |
| flag->count_lines | int  | Local  | 1     |
| flag->count_chars | char |        | 0     |

# Neck Properties

- Neck should be an articulation point in the CFG:
  - Dominator of main logic nodes
  - Always executed
    - Reachable from the entry
  - Executed once
    - Outside any loop structure

## Configuration Logic

```

struct Flags {
    char count_chars;
    int count_lines; };
int total_lines = 0;
int total_chars = 0;
int main(int argc, char** argv){
    struct Flags *flag;
    flag = malloc(sizeof(struct Flags));
    flag->count_chars = 0;
    flag->count_lines = 0;
    if (argc >= 2){
        for (int i = 1; i < argc; i++) {
            if (!strcmp(argv[i], "-c")) flag->count_chars = 1;
            if (!strcmp(argv[i], "-l")) flag->count_lines = 1; }}
  
```

**Neck** .....

```

char buffer[1024];
while (fgets(buffer, 1024, stdin)){
    if (count_chars) total_chars += sizeof(buffer);
    if (count_lines) total_lines++; }
if (count_chars) printf("# chars = %d\n", total_chars);
if (count_lines) printf("# Lines = %d\n", total_lines);
  
```

## Main Logic

# Partial Interpreter

---

Modified KLEE

# Program Partial State

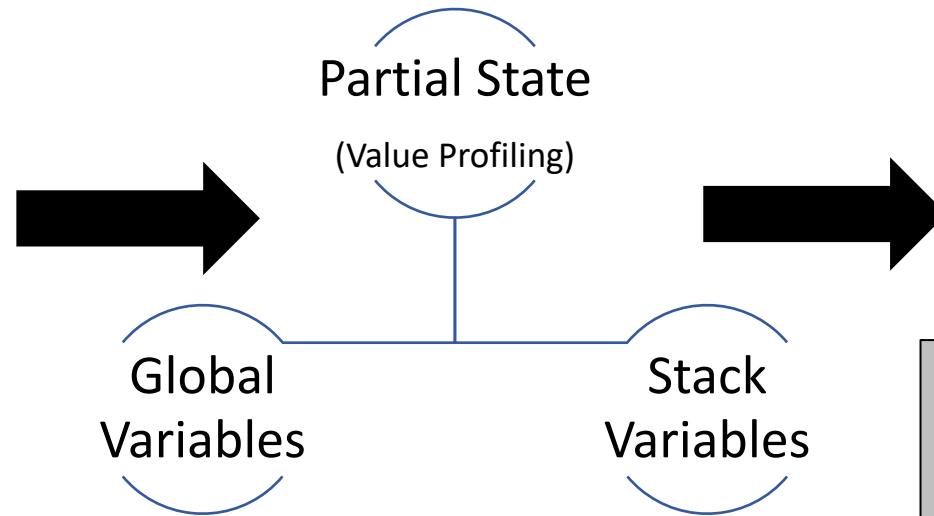
## Configuration Logic

```
struct Flags {
  char count_chars;
  int count_lines; };
int total_lines = 0;
int total_chars = 0;
int main(int argc, char** argv){
  struct Flags *flag;
  flag = malloc(sizeof(struct Flags));
  flag->count_chars = 0;
  flag->count_lines = 0;
  if (argc >= 2){
    for (int i = 1; i < argc; i++) {
      if (!strcmp(argv[i], "-c")) flag->count_chars = 1;
      if (!strcmp(argv[i], "-l")) flag->count_lines = 1; }}
```

Neck.....

```
char buffer[1024];
while (fgets(buffer, 1024, stdin)){
  if (count_chars) total_chars += sizeof(buffer);
  if (count_lines) total_lines++; }
if (count_chars) printf("# chars = %d\n", total_chars);
if (count_lines) printf("# Lines = %d\n", total_lines);
```

## Main Logic



| Variable          | Type | Scope  | Value |
|-------------------|------|--------|-------|
| total_lines       | int  | Global | 0     |
| total_chars       | int  | Global | 0     |
| flag->count_lines | int  | Local  | 1     |
| flag->count_chars | char |        | 0     |

### Stored as txt files:

- globals=gbls.txt
- primitive stack=primitiveLocals.txt
- struct stack=customizedLocals.txt
- ptr to struct=ptrToStructLocals.txt
- ptr to primitive=ptrToPrimitiveLocals.txt
- string Vars=stringVars.txt
- nested structs=nestedStructLocals.txt

# Partial Interpreter Implementation

**KLEE 2.1**

**LLVM 10.0**

# Compiler-Assisted Optimizations

---

Constant Conversion  
Multi-stage Simplifications



# Constant Conversion (CC)

## Configuration Logic

```
struct Flags {
  char count_chars;
  int count_lines; };
int total_lines = 0;
int total_chars = 0;
int main(int argc, char** argv){
  struct Flags *flag;
  flag = malloc(sizeof(struct Flags));
  flag->count_chars = 0;
  flag->count_lines = 0;
  if (argc >= 2){
    for (int i = 1; i < argc; i++) {
      if (!strcmp(argv[i], "-c")) flag->count_chars = 1;
      if (!strcmp(argv[i], "-l")) flag->count_lines = 1; }}
}
```



Backward CC

| Variable          | Type | Scope  | Value |
|-------------------|------|--------|-------|
| total_lines       | int  | Global | 0     |
| total_chars       | int  | Global | 0     |
| flag->count_lines | int  | Local  | 1     |
| flag->count_chars | char | Local  | 0     |



Forward CC

Neck.....

```
char buffer[1024];
while (fgets(buffer, 1024, stdin)){
  if (count_chars) total_chars += sizeof(buffer);
  if (count_lines) total_lines++; }
if (count_chars) printf("# chars = %d\n", total_chars);
if (count_lines) printf("# Lines = %d\n", total_lines);
```

## Main Logic

## Constant Conversion LLVM Pass

- globals=**gbls.txt**
- plocals=**primitiveLocals.txt**
- clocals=**customizedLocals.txt**
- ptrStructlocals=**ptrToStructLocals.txt**
- ptrToPrimLocals=**ptrToPrimitiveLocals.txt**
- stringVars=**stringVars.txt**
- nestedStrcts=**nestedStructLocals.txt**

# CC Pre-neck

```
1 int main (int argc, char **argv)
2 {
3     struct rm_options x;
4     rm_option_init (&x);
5
6     while ((c = getopt_long (argc, argv, "dfirvIR", long_opts, NULL)) != -1)
7     {
8         switch (c)
9         {
10            {
11                case 'd':
12                    x.remove_empty_directories = true;
13                    break;
14
15                case 'f':
16                    x.interactive = RMI_NEVER;
17                    x.ignore_missing_files = true;
18                    prompt_once = false;
19                    break;
20
21                case 'i':
22                    x.interactive = RMI_ALWAYS;
23                    x.ignore_missing_files = false;
24                    prompt_once = false;
25                    break;
26            }
27        }
28    }
```

rm GNU Coreutils

```
1 static void rm_option_init (struct rm_options *x)
2 {
3     x->ignore_missing_files = false;
4     x->interactive = RMI_SOMETIMES;
5     x->one_file_system = false;
6     x->remove_empty_directories = false;
7     x->recursive = false;
8     x->root_dev_ino = NULL;
9     x->preserve_all_root = false;
10    x->stdin_tty = isatty (STDIN_FILENO);
11    x->verbose = false;
12    x->require_restore_cwd = false;
13 }
```

Specialize API  
rm\_option\_init

# Multi-stage Simplifications


```

b := 3
c := 1 + b
d := b + c
  
```

→

```

b := 3
c := 1 + 3
d := 3 + c
  
```



## Constant Propagation

- Standard LLVM pass

## Simplifying CFG

- Standard LLVM Pass



## Cleaning up

- Customized LLVM Pass

# Cleaning up

## Removing unused functions

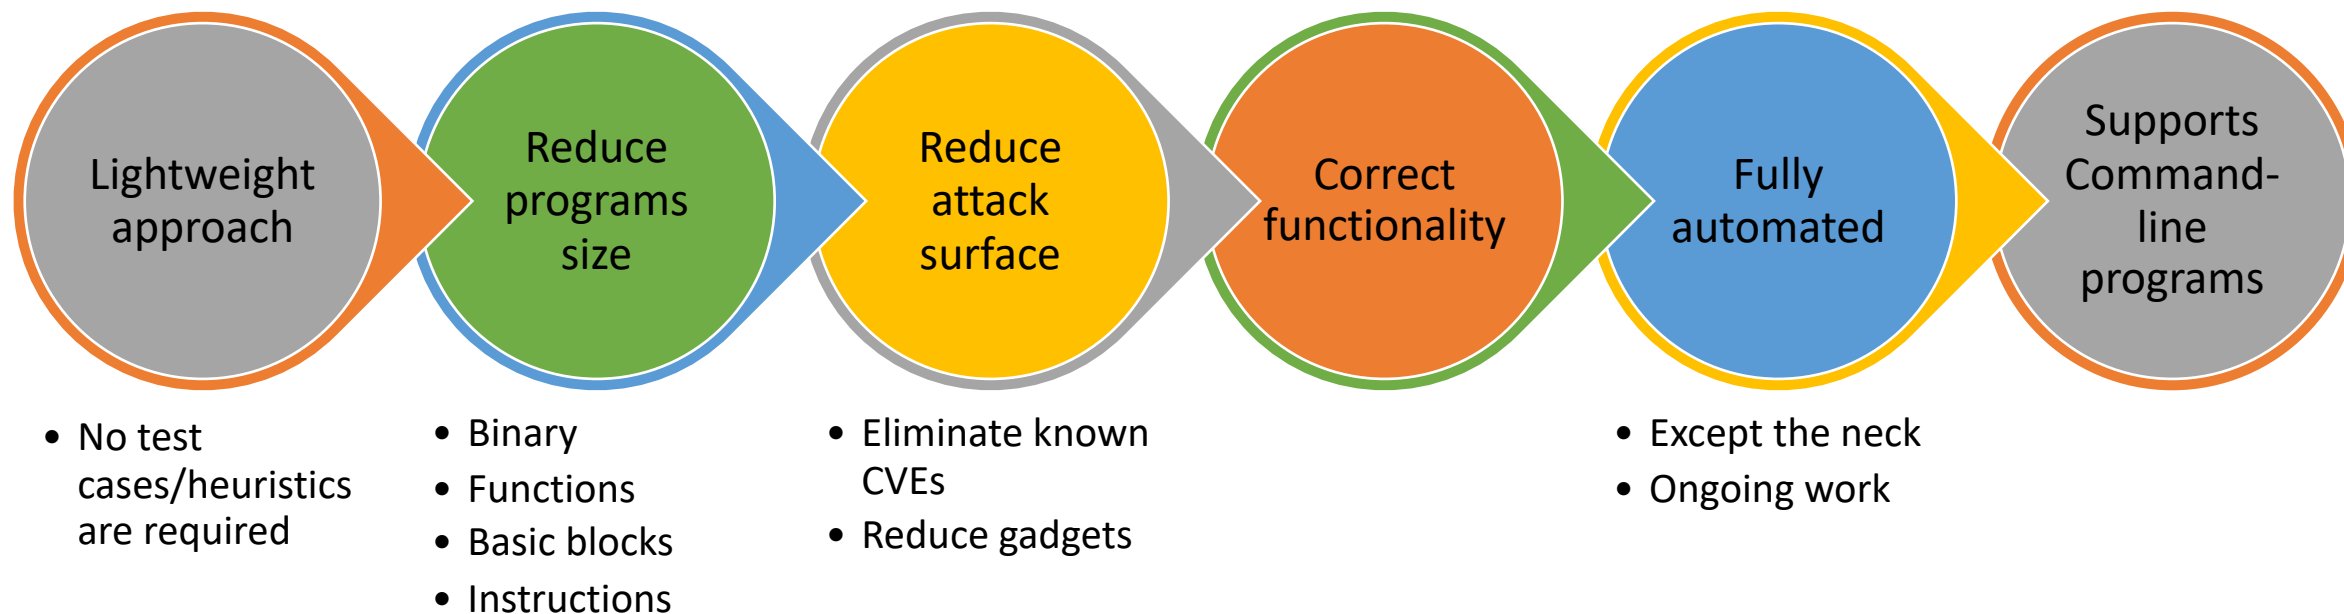
- Iterate the Call Graph
- Check the number of users of each function
- Remove the function if number of users is zero
  - function pointer won't be removed

## Removing unused variables in the remained functions

- Local and global

We used users() LLVM API

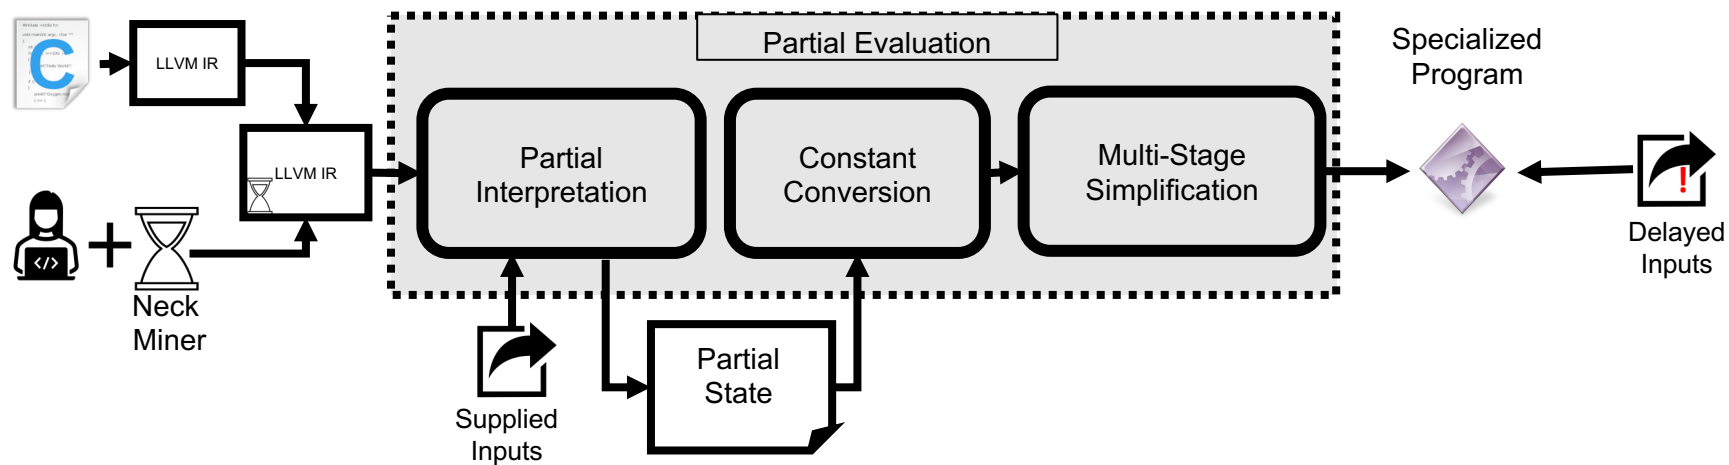
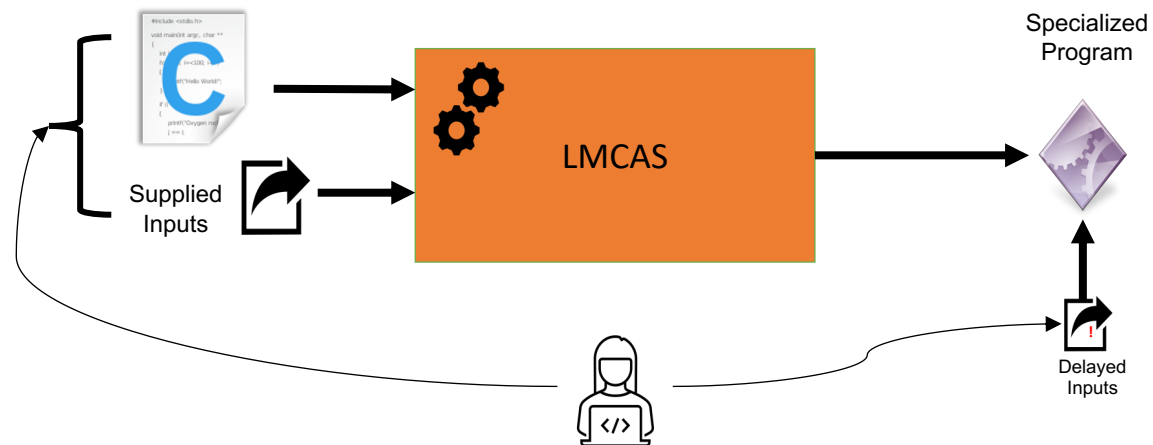
# LMCAS Power



# Agenda

- Use Case
- Introducing LMCAS
- **Questions**
- Demo

# LMCAS - Summary



# Agenda

- Use Case
- Introducing LMCAS
- Questions
- **Demo**





# Demo Roadmap

---

Two debloating scenarios

---

See size reduction statistics

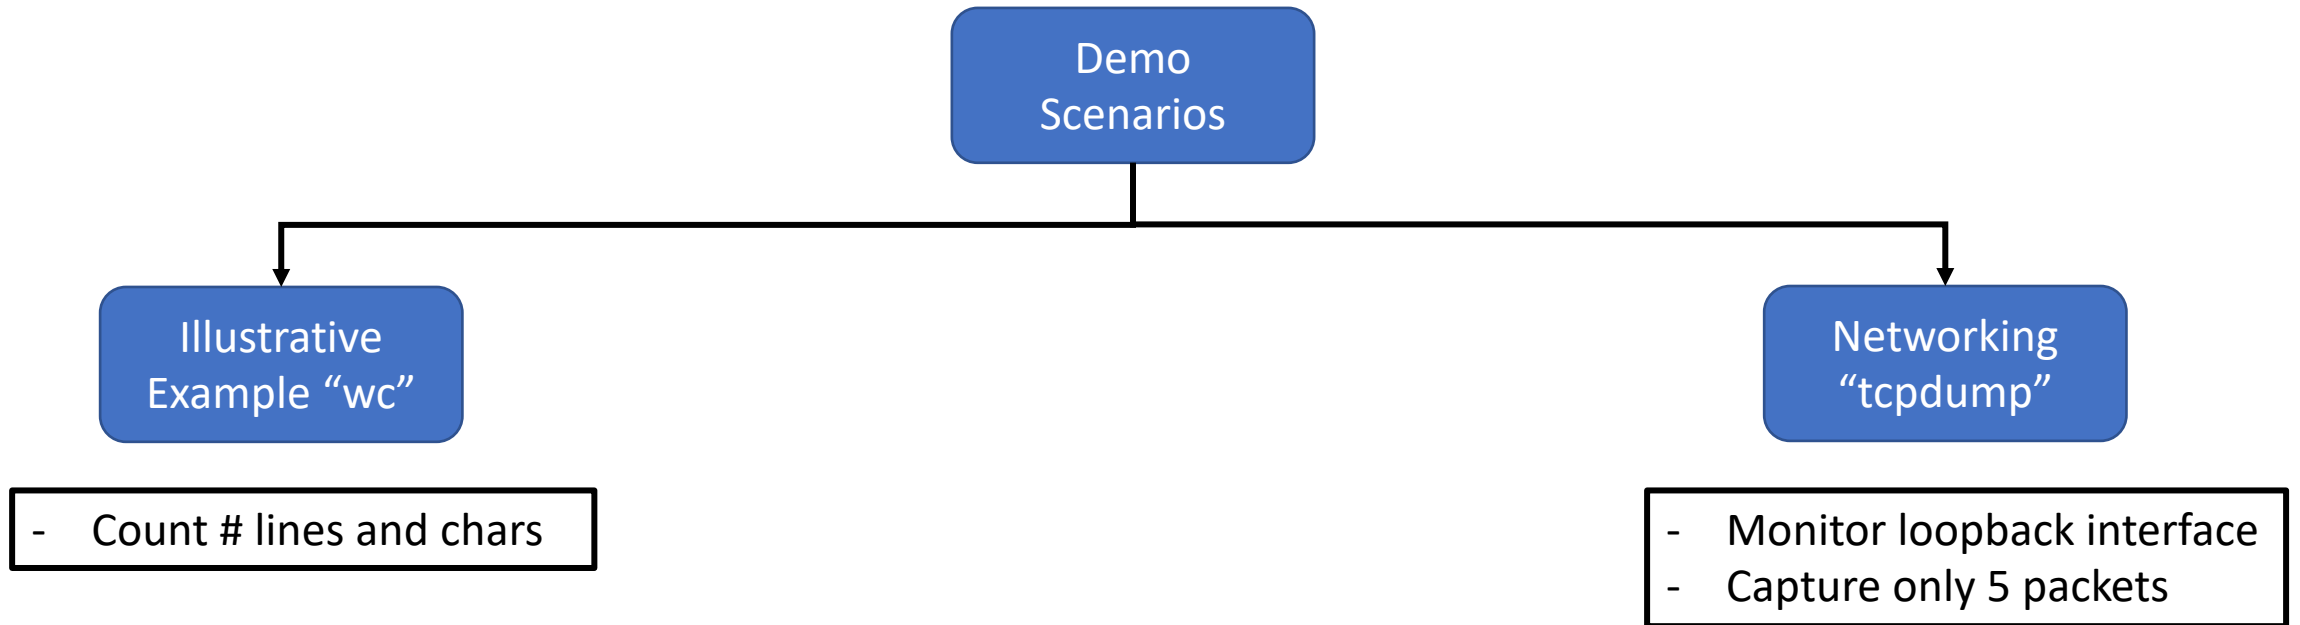
---

See gadgets reduction statistics

---

Verify the functionality of the debloated program

# Debloating Scenarios



LMCAS Docker and programs used in the demo are available at:  
[https://github.com/Mohannadcse/LMCAS\\_Demo](https://github.com/Mohannadcse/LMCAS_Demo)

**18 Programs**

# LMCAS Demo Repo

main ▾ LMCAS\_Demo / LMCAS\_Docker /

Mohannadcse remove duplicated directory

..

- bitcode\_files
- build**
- source\_code\_files
- .gitignore
- Dockerfile
- create\_histograms.py
- ropAnalysis.sh
- runDemo.sh**
- statistics.sh
- testScript.sh

## WLLVM

main ▾ LMCAS\_Demo / LMCAS\_Docker / bitcode\_files /

Mohannadcse renaming

..

|                  |          |
|------------------|----------|
| basename_orig.bc | renaming |
| basenc_orig.bc   | renaming |
| comm_orig.bc     | renaming |
| date_orig.bc     | renaming |
| du_orig.bc       | renaming |
| echo_orig.bc     | renaming |
| fmt_orig.bc      | renaming |
| fold_orig.bc     | renaming |
| head_orig.bc     | renaming |
| id_orig.bc       | renaming |
| kill_orig.bc     | renaming |
| objdump.bc       | renaming |
| readelf.bc       | renaming |
| realpath_orig.bc | renaming |
| sort_orig.bc     | renaming |
| tcpdump.bc       | renaming |
| uniq_orig.bc     | renaming |
| wc_orig.bc       | renaming |

## klee\_dump\_memory


main ▾ LMCAS\_Demo / LMCAS\_Docker / source\_code\_files /

Mohannadcse renaming

..

|            |          |
|------------|----------|
| basename.c | renaming |
| basenc.c   | renaming |
| comm.c     | renaming |
| date.c     | renaming |
| du.c       | renaming |
| echo.c     | renaming |
| fmt.c      | renaming |
| fold.c     | renaming |
| head.c     | renaming |
| id.c       | renaming |
| kill.c     | renaming |
| objdump.c  | renaming |
| readelf.c  | renaming |
| realpath.c | renaming |
| sort.c     | renaming |
| tcpdump.c  | renaming |
| uniq.c     | renaming |
| wc.c       | renaming |

# Specializing GNU wc



# Specializing tcpdump

mp

# Summary



LMCAS specializes programs

Compiler Optimizations  
Disciplined Implementation



LMCAS Benefits

Lightweight approach  
Reducing attack surface  
Reducing program size



Ongoing project

Automating the neck identification  
Leveraging data flow analysis  
Covering more real-world use cases

# More Info



[https://github.com/Mohannadcse/LMCAS\\_Demo](https://github.com/Mohannadcse/LMCAS_Demo)

**arXiv**

<https://arxiv.org/abs/2109.02775>



[mohannad@cs.wisc.edu](mailto:mohannad@cs.wisc.edu)



Malhanahnah