Programming Languages Research Overview

Stefan Muller

CS100

Oct. 8, 2021

Short Bio

- PhD in Computer Science from Carnegie Mellon in 2018
- Postdoc at Carnegie Mellon, 2018-2020
- Joined IIT in 2020 as Assistant Professor of Computer Science
- Research interests:
  - Language and type system design
  - Static resource analysis
  - Parallel computing

We're more and more dependent on software

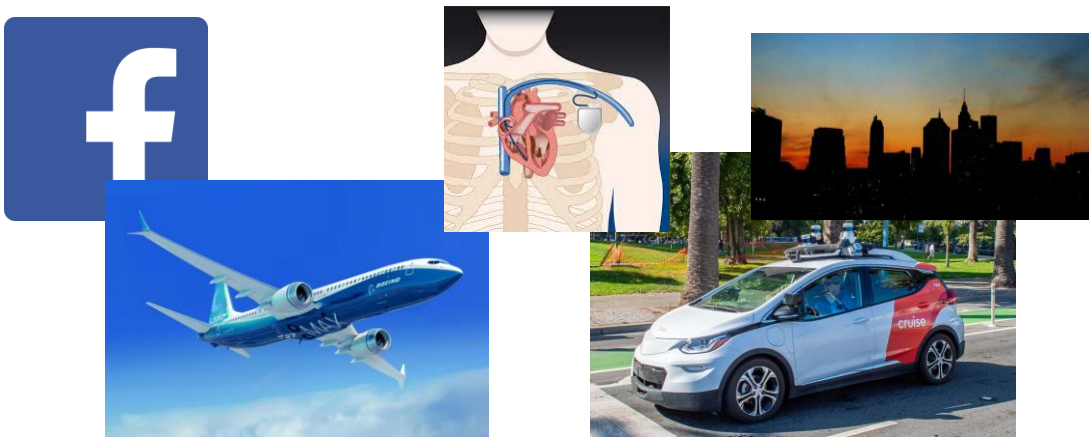

3

… and people are bad at writing software

especially

large
concurrent
interactive
time-sensitive
multi-language
…

software

4

Even getting simple programs right is hard

```
sort([8;2;1;6;3]) = [1;2;3;6;8]

def sort (l) … ?
```

5

It would be great if the language gave us some help!

- Python `def sort (l)`
  - Takes… something and returns… something
  - Valid: `sort([8;2;1;6;3]) = "Hello World"`
- C `node *sort(node *list)`
  - Takes a pointer to a node and returns a pointer to a node.
  - Valid: `sort([8;2;1;6;3]) = ['H'; 'e'; 'l'; 'l'; 'o'; …]`

6

It would be great if the language gave us some help!

- OCaml `sort : int list -> int list`
  - Takes an integer list and returns an integer list.
  - Valid: `sort([8;2;1;6;3]) = [8;2;1;6;3]`
  - Valid: `sort([8;2;1;6;3]) = [10;11;12]`
- Coq
```
sort : forall (l1 : list int), exists (l2: list int),
                    Sorted l2 /\ Permutation l1 l2
```
  - Takes an integer list and returns a sorted permutation of it.
  - Valid: `sort([8;2;1;6;3]) = [1;2;3;6;8]`
  - … and nothing else

7

Concurrency adds more complexity

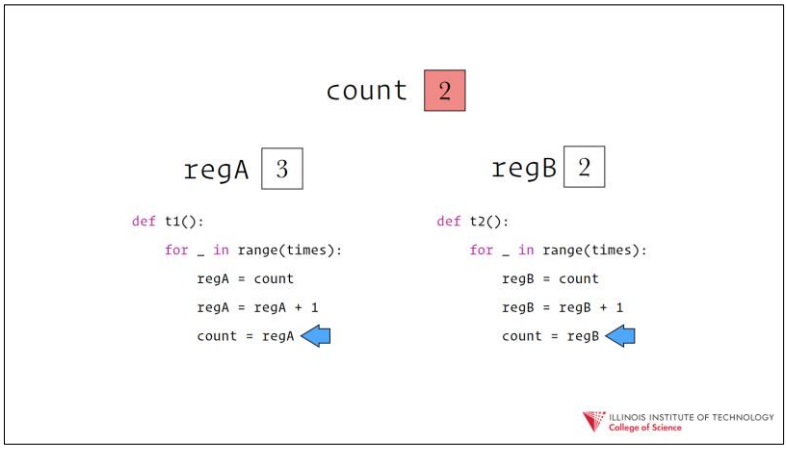

A race condition!

CS100, 9/17

8

Minimizing shared state helps

- Many languages based around highly parallel, no-shared-state applications

```
parallel for i in range(1000000):
  A[i] = B[i] + C[i]
```

0-999 1000-1999 2000-2999….

9

Minimizing shared state helps … to a point

```
parallel for i in range(1000000):
  if (i == 0):
     name = input(“What is your name?”)
     print(“Hello, ” + name)
  else:
    A[i] = B[i] + C[i]
```

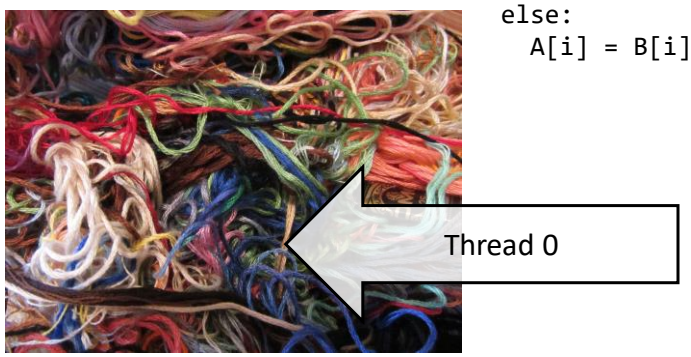

10

Solution: thread priorities

```
priority High
priority Low
order Low < High

parallel at High:
  name = input("What is your name?")
  print("Hello, " + name)
parallel at Low:
  parallel for i in range(1, 1000000):
    A[i] = B[i] + C[i]
```

11

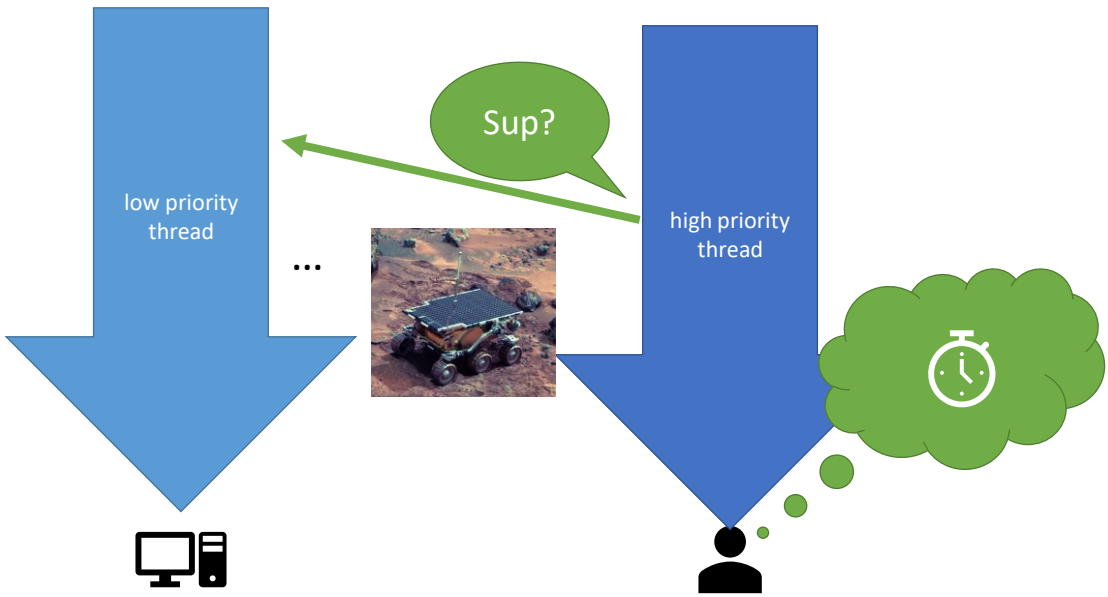

12

It would be great if the language gave us some help!

- PriML (PhD Dissertation, ongoing work)
  - Prevent priority inversions **at compile time**

```
constraint violated at example.prm:5.1-5.8 :  high  <=  low
Type error: constraint violated
```

13

Current research projects

- PriML: parallel interactive programs
- RaCUDA: Resource-aware CUDA
  - How long will your GPU kernel take to run?
- SEEr: Scalable, Energy-Efficient HPC environment for AI-enabled science
  - (with Zhiling)

14

RaCUDA: Resource-aware CUDA

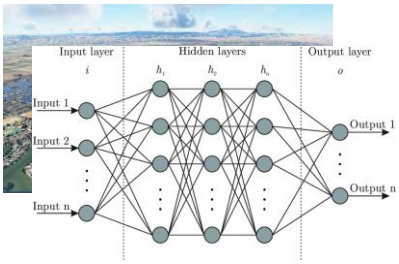

GPGPU: General-purpose programming for GPUs
CUDA: Extension of C for GPGPU programming

15

Writing efficient CUDA programs is hard

- CUDA (Idea):

C programmer

- CUDA (Practice)

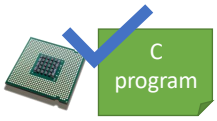

16

Writing efficient CUDA programs is hard

- RaCUDA (Current):

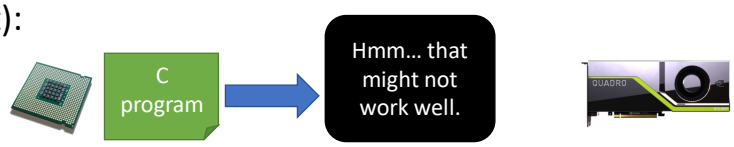

- RaCUDA (Soon, hopefully!)

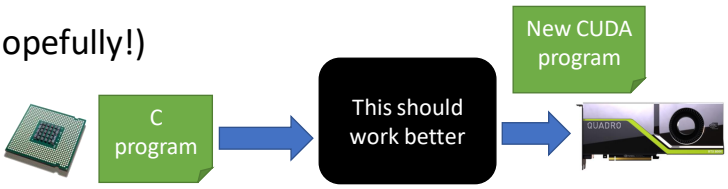

17

Current research projects

- PriML: parallel interactive programs
- RaCUDA: Resource-aware CUDA
  - How long will your GPU kernel take to run?
- SEEr: Scalable, Energy-Efficient HPC environment for AI-enabled science
  - (with Zhiling)

18

Running software on large systems requires making lots of decisions

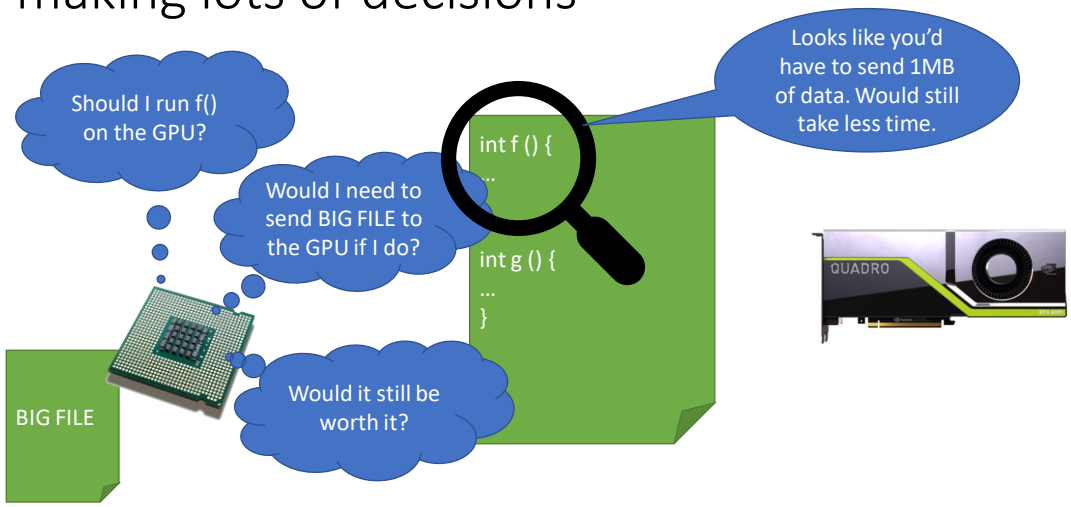

19

Collaborators

- PriML: parallel interactive programs


Umut Acar (CMU)


Kunal Agrawal
Wash U. in St. Louis


Angelina Lee

- RaCUDA: Resource-aware CUDA


Jan Hoffmann (CMU)

- SEEr: Scalable, Energy-Efficient HPC environment for AI-enabled science


Zhiling Lan
IIT


Romit Maulik
IIT/Argonne


Mike Papka
NIU/Argonne


Valerie Taylor
U Chicago/Argonne


Xingfu Wu

20

Students

- Xiangwei (Shawn) Li (Masters)
- Mark Lou (Undergrad)
- Deepika Padmanabhan (Masters)

- Research Opportunities Available!
  - Helpful background:
    - PL (e.g., CS440)
    - Functional programming (OCaml, Haskell, etc., CS 340)

- Contact: smuller2@iit http://cs.iit.edu/~smuller/

Current research projects

- PriML: parallel interactive programs
- RaCUDA: Resource-aware CUDA
- SEEr: Scalable, Energy-Efficient HPC environment for AI-enabled science

21