

# SEED: Confidential Big Data Workflow Scheduling with Intel SGX Under Deadline Constraints

Ishtiaq Ahmed<sup>†</sup>, Saeid Mofrad<sup>†</sup>, Shiyong Lu<sup>†</sup>, Changxin Bai<sup>†</sup>, Fengwei Zhang<sup>★</sup>,  
Dunren Che<sup>★</sup>

Wayne State University, USA<sup>†</sup>

Southern University of Science and Technology, China<sup>★</sup>

Southern Illinois University, USA<sup>★</sup>



DEPARTMENT OF COMPUTER SCIENCE

World-Class Education in the Real World™



# Overview

- ▶ Introduction
- ▶ The Problem
- ▶ Existing Approaches and Limitations
- ▶ Our Proposed Approach
- ▶ Experimental Results
- ▶ Conclusions and future work

# Introduction

- ▶ Big data workflow management systems [2] (BDWFMSs) have recently emerged as popular data analytics platforms to perform large-scale data analytics in the cloud.
- ▶ The protection of data confidentiality and secure execution of workflow applications remains an important and challenging problem.
- ▶ Monetary cost optimization of executing workflows in the cloud while satisfying a user-defined deadline is also challenging.

# The problem of big data workflow scheduling in the cloud

- ▶ The Protection of the privacy of the confidential workflow's tasks in the cloud, whose proprietary algorithm implementations are intellectual properties of the respective stakeholders.
- ▶ The monetary cost optimization of executing workflows in the cloud while satisfying a user-defined deadline.

# Limitation of previous approaches

- ▶ Require many VM instances.
- ▶ Operational cost is expensive.
- ▶ Scheduling algorithm that can handle confidential tasks is scarce.



What if I need low operational cost with confidentiality and integrity environment?

# Challenges for scheduling algorithms

- ▶ **Deadline–centric:** When the deadline is fixed, how we schedule a workflow with minimal cost?
- ▶ **Budget–centric:** When the budget is fixed how we schedule a workflow with minimal makespan?

# Contributions

- ▶ We propose a novel deadline-constrained and SGX-aware big data workflow scheduling algorithm, SEED (SGX, Efficient, Effective, Deadline Constrained) that leverage hardware-based security features such as Intel Software Guard eXtensions (SGX) [1] in the cloud hardware to protect the execution of security-sensitive workflow's tasks in the cloud environments.
- ▶ SEED features several heuristics, including exploiting the longest critical paths and reuse of extra times in existing VMs.
- ▶ The proposed SEED, as well as IC-PCP, have been tested in a modified version of DATAVIEW [2], which is one of the most usable big data workflow management systems in the community, to support confidential tasks in a workflow DAG with Intel SGX.

# Planner level: Scheduling workflow

---

**Algorithm 1: ScheduleWorkflow( $G$ )**

---

- 1  $startNodes \leftarrow nodesWithoutParent$  in  $G$ ;
  - 2  $endNodes \leftarrow nodesWithoutChildren$  in  $G$ ;
  - 3 connect  $t_{start}$  with  $startNodes$ ;
  - 4 connect  $t_{end}$  with  $endNodes$ ;
  - 5 determine the available virtual machines with cost and billing unit cycle;
  - 6  $EST(t_{start}) \leftarrow 0$ ;
  - 7  $LFT(t_{end}) \leftarrow \text{Deadline } \delta$ ;
  - 8 calculate  $EST, EFT$  from  $t_{start}$  to  $t_{end}$  recursively;
  - 9 calculate  $LFT$  from  $t_{end}$  to  $t_{start}$  recursively;
  - 10  $G \leftarrow$  remove  $t_{start}, t_{end}$ , and edges in between from  $G$ ;
  - 11  $AssignChildren(G)$ ;
-

# Planner level: Assigning Children

---

**Algorithm 2:** AssignChildren( $G$ )

---

```
1  $list \leftarrow$  sort tasks in  $G$  in topological order;
2 while  $list \neq \emptyset$  do
3    $criticalPath \leftarrow \emptyset$ ;
4    $t_i \leftarrow list[1]$  ▷ Selecting first task
5   add  $t_i$  in  $criticalPath$ ;
6   while  $t_i$  has children do
7      $t_i \leftarrow criticalChild(t_i)$ ;
8     add  $t_i$  in  $criticalPath$ ;
9   end
10   $paths \leftarrow$  split apart  $criticalPath$  by confidential tasks as
    separators;
11  foreach  $path \in paths$  do
12     $AllocateVirtualMachine(path)$ ;
13    foreach  $t_i \in path$  do
14      recalculate LFT of the ancestors  $\{t_k : t_k \notin path\}$  of
         $t_i$ ;
15      recalculate EST, EFT of the descendants
         $\{t_k : t_k \notin path\}$  of  $t_i$ ;
16    end
17     $G \leftarrow$  remove all tasks and edges involving tasks in  $path$ 
        from  $G$ ;
18     $list \leftarrow$  remove all tasks involving in  $path$  from  $list$ ;
19  end
20 end
```

---

# Planner level: Allocating virtual machines

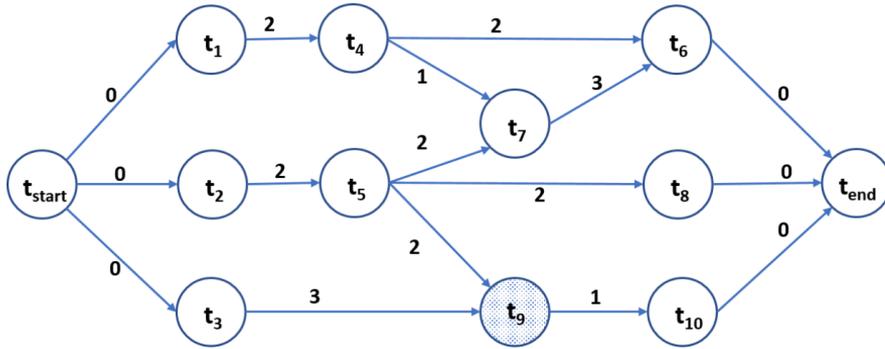
---

**Algorithm 3:** AllocateVirtualMachine(*path*)

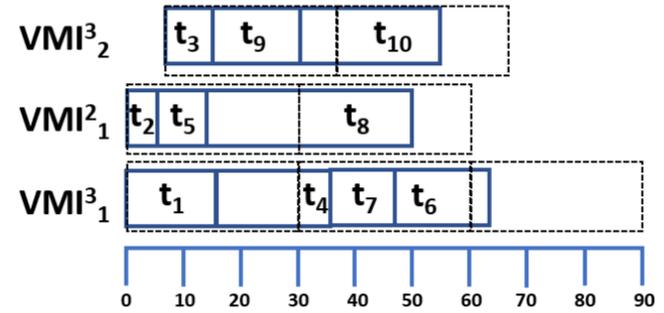
---

```
1  $VMI^c \leftarrow$  existing-VM-instances;
2 if secure environment is required for path then
3   launch a new instance  $vmi^{SGX}$  for path;
4    $VMI^c \leftarrow vmi^{SGX}$ ;
5   return;
6 foreach existing VM instance  $vmi \in VMI^c$  in ascending order of
   cost do
7   oldAssignment  $\leftarrow$  assignment(vmi);
8    $Q \leftarrow$  assignment(vmi);  $\triangleright$  queued tasks at vmi
9    $\triangleright$  Case 1: Insert path into middle of  $Q$ 
10  if there exists a task  $t_k \in Q$  such that  $t_k$ 
11   $\in path[last].children$  then
12    insert path into  $Q$  before  $t_k$  and recalculate EST, EFT,
13    LFT of affected tasks in  $Q$ ;
14    if assignment(vmi) is valid then
15      return;  $\triangleright$  keep current assignment(vmi)
16    else
17      assignment(vmi)  $\leftarrow$  oldAssignment;
18    end
19   $\triangleright$  Case 2: Insert path at head of  $Q$ 
20  insert path at head  $Q$  and recalculate EST, EFT, LFT of
21  affected tasks in  $Q$ ;
22  if assignment(vmi) is valid then
23    return;  $\triangleright$  keep current assignment(vmi)
24  else
25    assignment(vmi)  $\leftarrow$  oldAssignment;
26  end
27   $\triangleright$  Case 3: Insert path at end of  $Q$ 
28  insert path at end of  $Q$  and recalculate EST, EFT, LFT of
29  affected tasks in  $Q$ ;
30  if assignment(vmi) is valid then
31    return;  $\triangleright$  keep current assignment(vmi)
32  else
33    assignment(vmi)  $\leftarrow$  oldAssignment;
34  end
35 end
36  $\triangleright$  No existing VM instance can accommodate path
37 launch a new cheapest but capable VM instance  $vmi^c$  for executing
38 the tasks in path before their LFT;
39  $VMI^c \leftarrow vmi^c$ ;
40 return
```

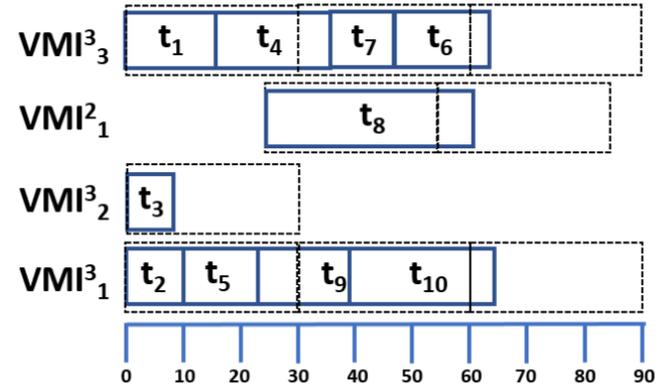
---



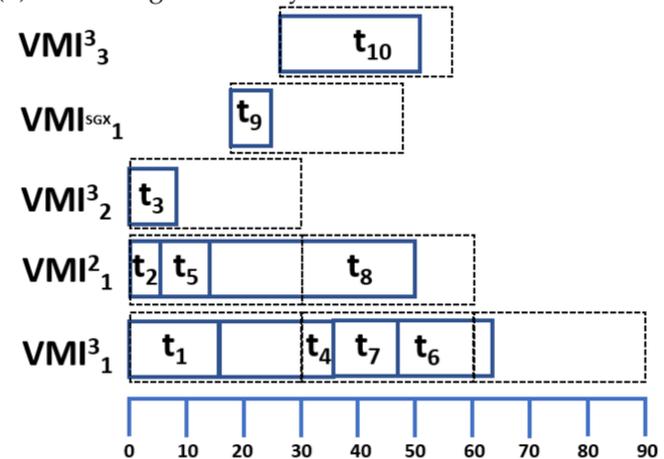
Type	Start	End	Duration	Interval cycle	Cost	Selected task
$VMI_1^3$	0	66	66	3	15	$\{t_1, t_4, t_7, t_6\}$
$VMI_1^2$	0	51	51	2	16	$\{t_2, t_5, t_8\}$
$VMI_2^3$	8	58	50	2	10	$\{t_3, t_9, t_{10}\}$



(a) Schedule generated by SGX-E2C2D on non-confidential workflow.



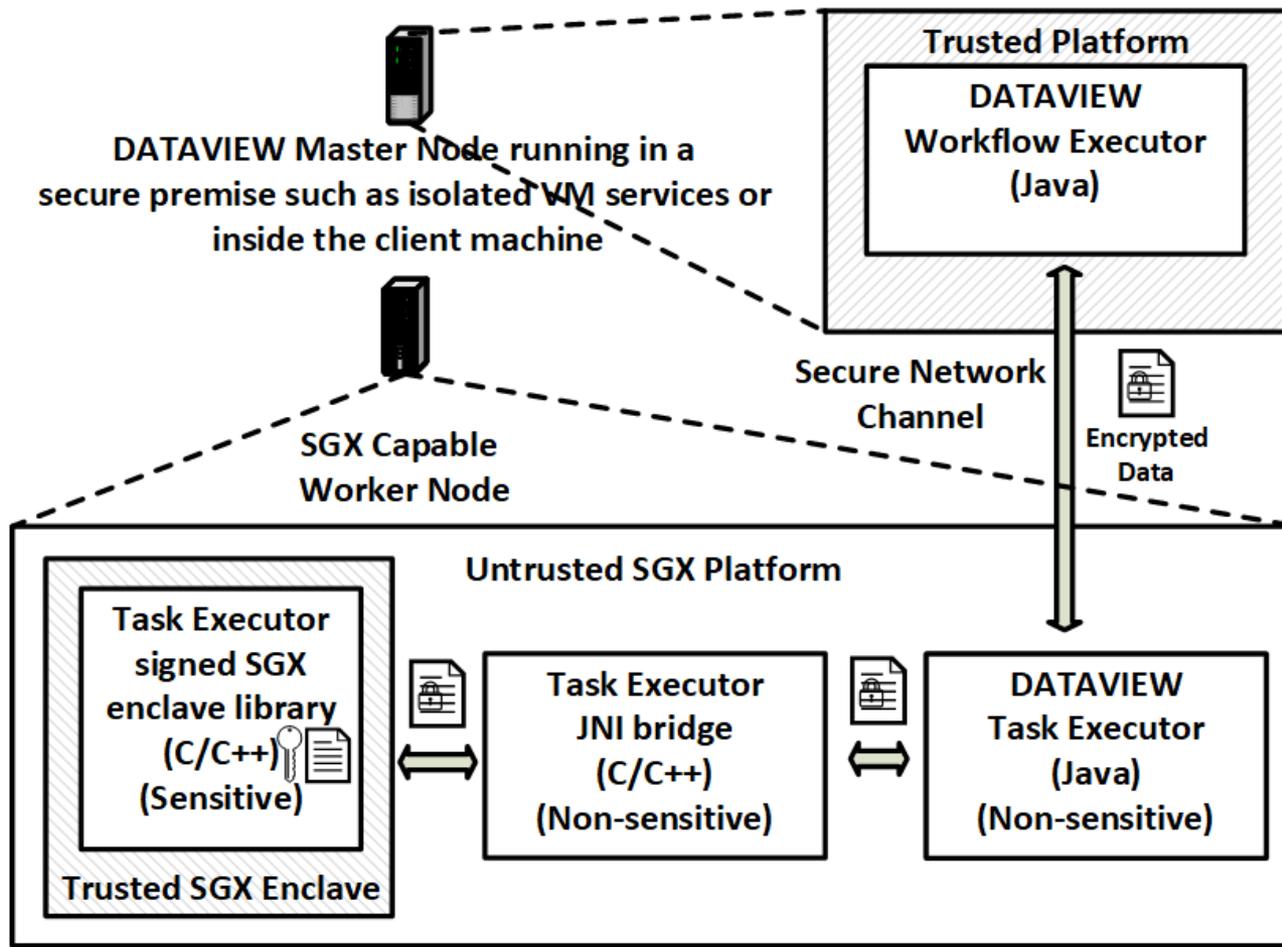
(b) Schedule generated by IC-PCP on non-confidential workflow.



(c) Schedule generated by SGX-E2C2D on confidential workflow.

Type	Start	End	Duration	Interval cycle	Cost	Selected task
$VMI_1^3$	0	66	66	3	15	$\{t_1, t_4, t_7, t_6\}$
$VMI_1^2$	0	51	51	2	16	$\{t_2, t_5, t_8\}$
$VMI_2^3$	0	9	9	1	5	$\{t_3\}$
$VMI_{SGX}^3$	17	24	7	1	25	$\{t_9\}$
$VMI_3^3$	25	50	25	1	5	$\{t_{10}\}$

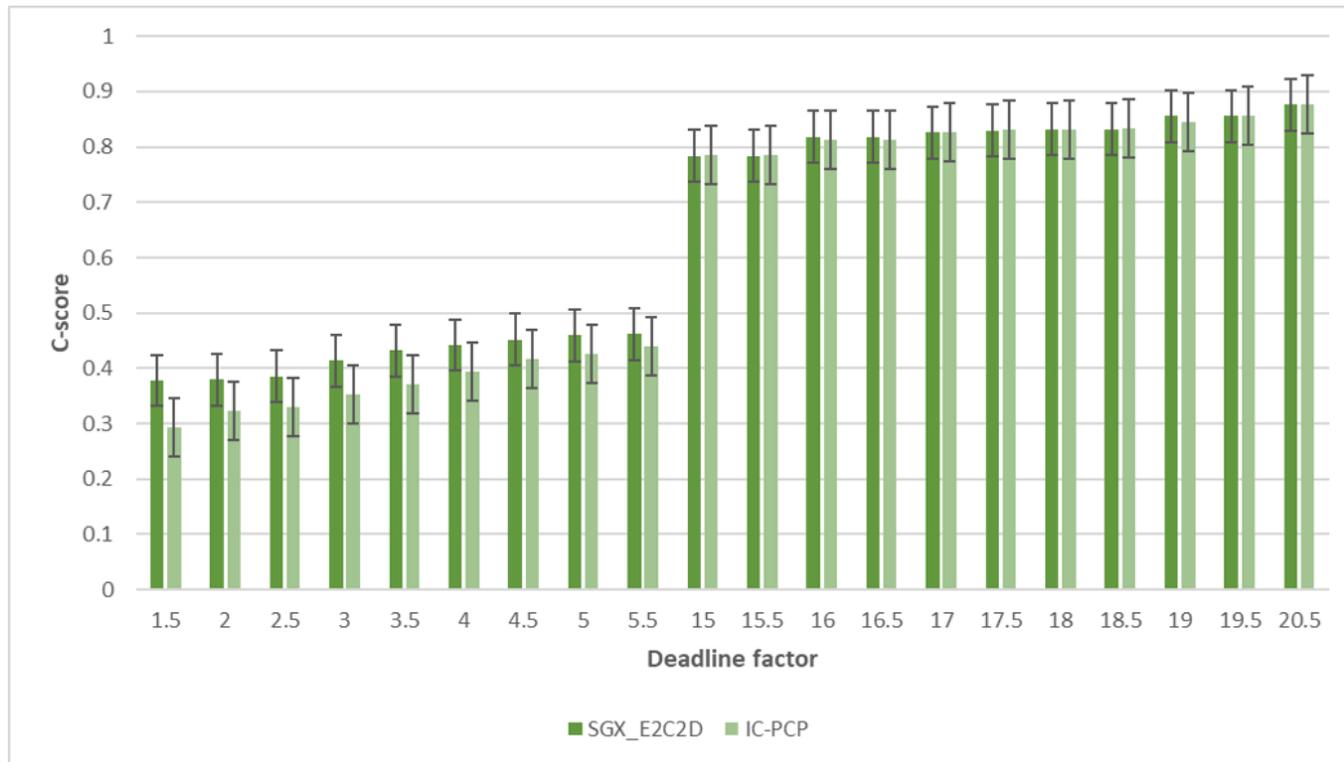
# The architecture of DATAVIEW Task Executor with SGX support



# Experimental settings

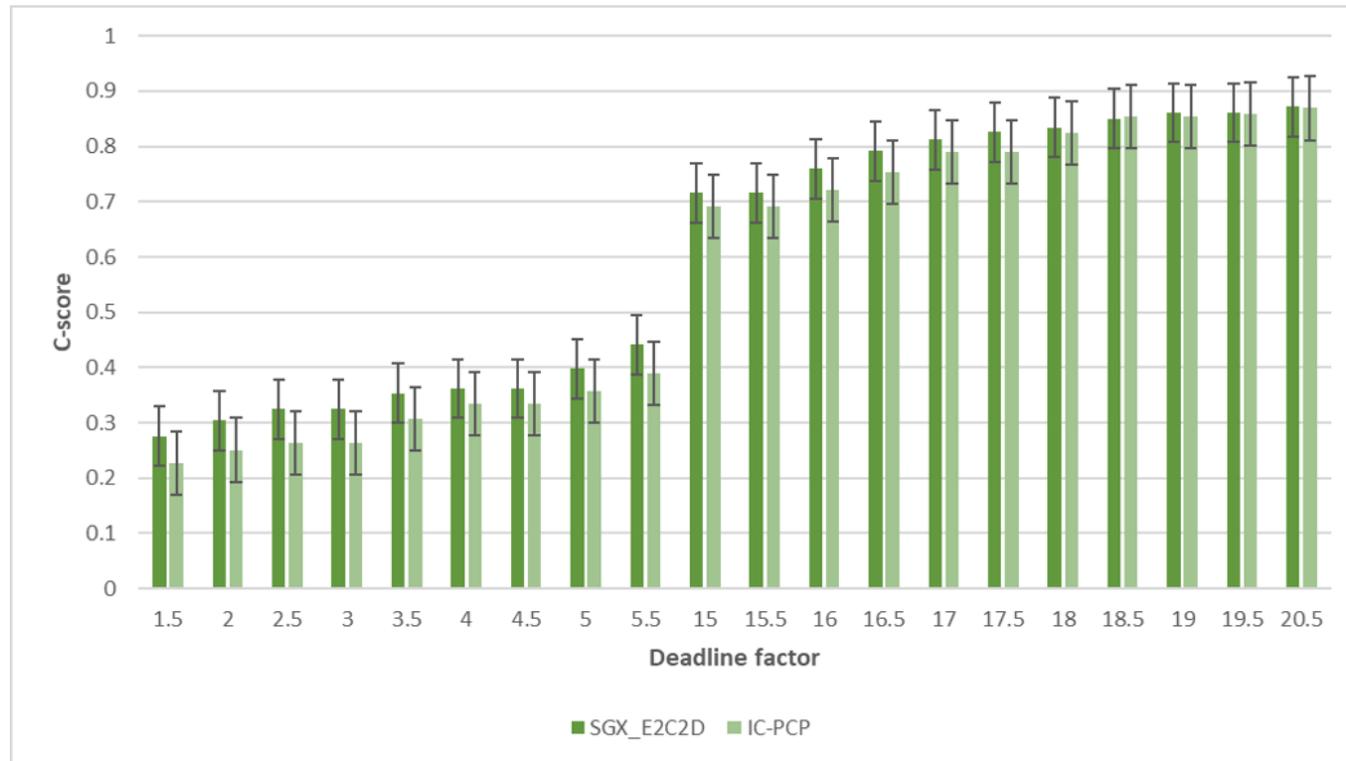
<b>Testbed Machine</b>	<b>Intel SGX</b>	<b>AWS</b>
<b>CPU Model</b>	Xeon E3-1275 V6	Intel(R) Xeon(R) CPU E5-2676 v3
<b>CPU Physical Core Number</b>	4	8
<b>CPU Logical Thread Number</b>	8	16
<b>CPU Base Clock</b>	3.8GHz	2.4GHz
<b>CPU Boost Clock</b>	4.2GHz	3.2GHz
<b>Cache Type</b>	Smart Cache	Smart Cache
<b>Cache Size</b>	8MB	20MB
<b>Motherboard</b>	Supermicro X11SSW-F	AWS
<b>Memory</b>	32GB DDR4 ECC	1GB DDR4 ECC
<b>Storage</b>	HDD	EBS
<b>Operating System</b>	CentOS 7.0	Linux 16.04 LTS
<b>Kernel Version</b>	3.10.0-862.9.1.el7.x86_64 x86_64	4.4.0-1060-aws
<b>SGX SDK Version</b>	SGX SDK Ver 2.2	SGX SDK Ver 2.0
<b>AMI family</b>	N/A	General purpose
<b>AMI type</b>	N/A	t2.micro
<b>AMI Assigned CPU Core</b>	N/A	1

# Performance analysis Montage Workflow (10sec)

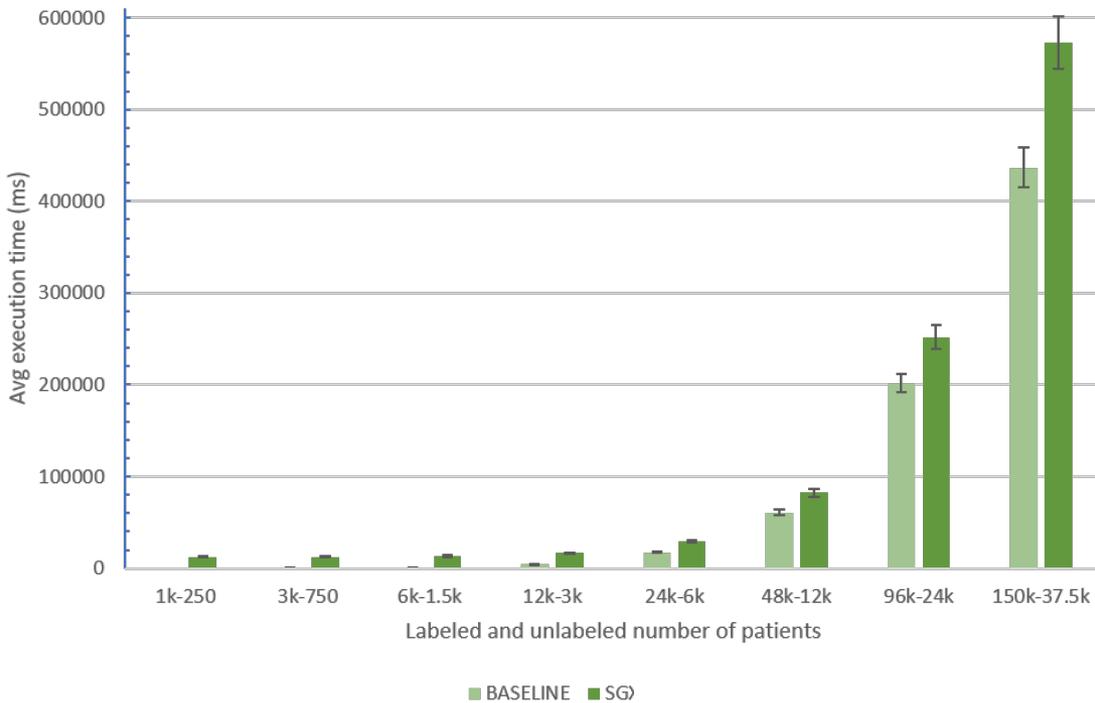


$$C = \begin{cases} 0.5 + 0.5 * \frac{(maxcost - cost)}{maxcost}, & \text{if } makespan \leq \delta \\ 0.5 - 0.5 * \frac{(makespan - \delta)}{(maxmakespan - \delta)}, & \text{otherwise} \end{cases}$$

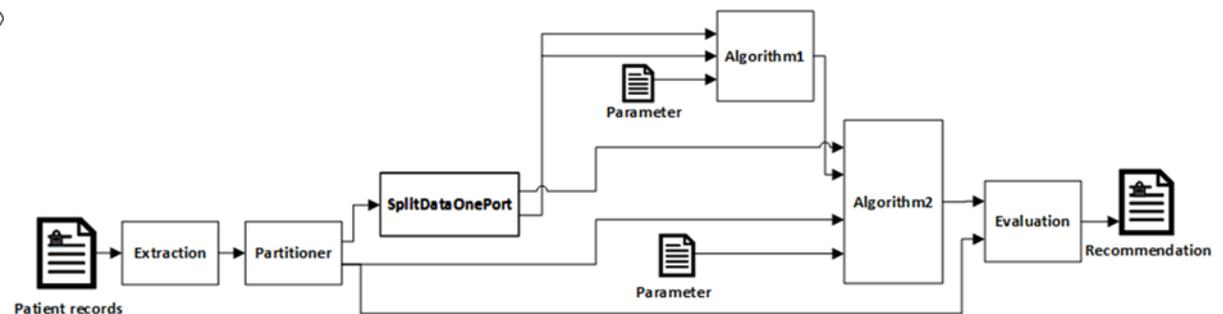
# Performance analysis Montage Workflow (60s)



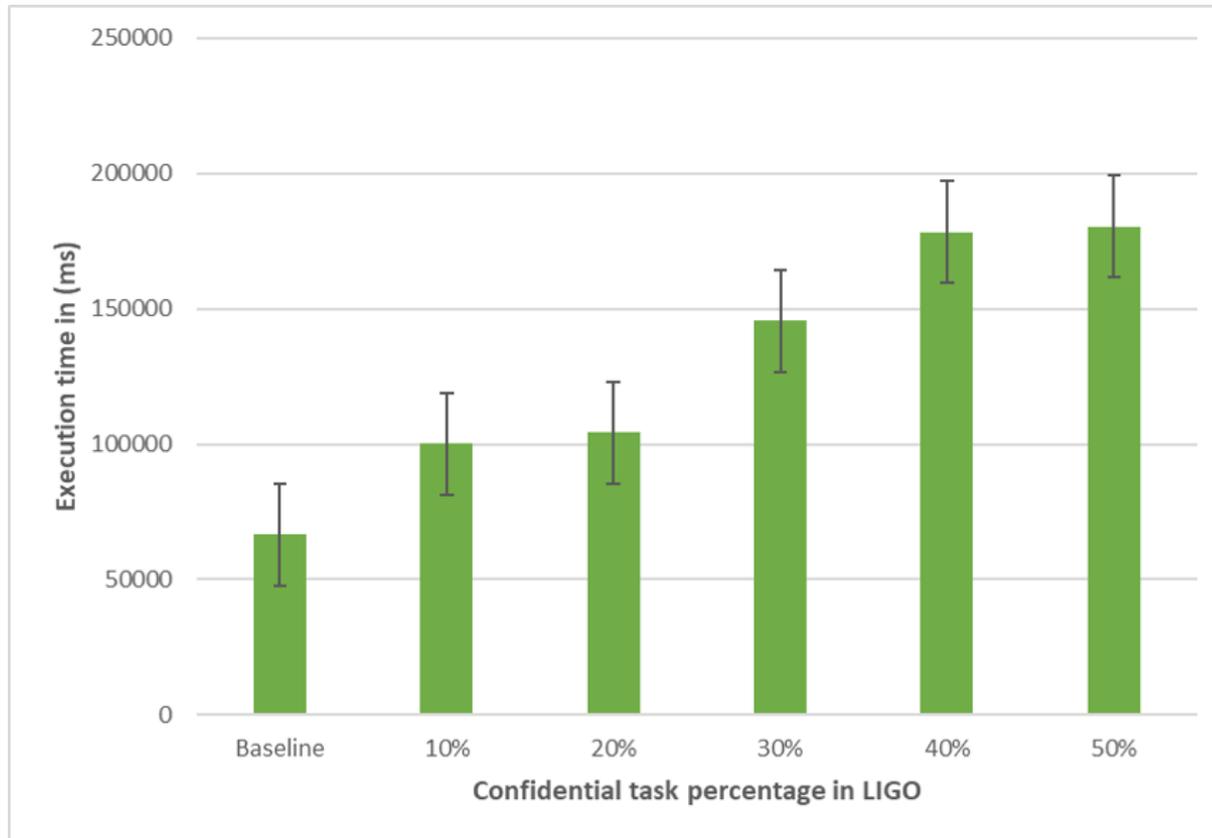
# Performance analysis of algorithm1 in the diagnosis recommendation workflow[3]



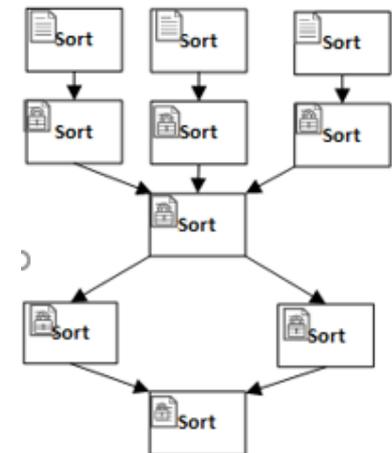
- SGX overhead reported a range of 1.24x to 1.31x overhead compared with the base line with 96k–150k labeled and 24k–37.5k unlabeled instances, respectively



# Performance overhead of LIGO workflow



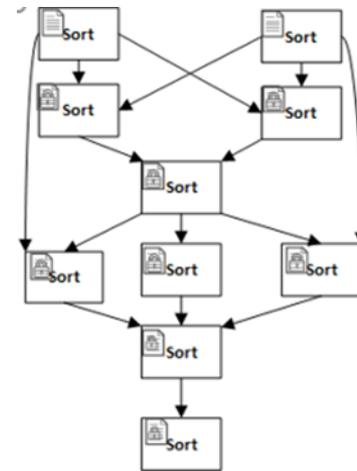
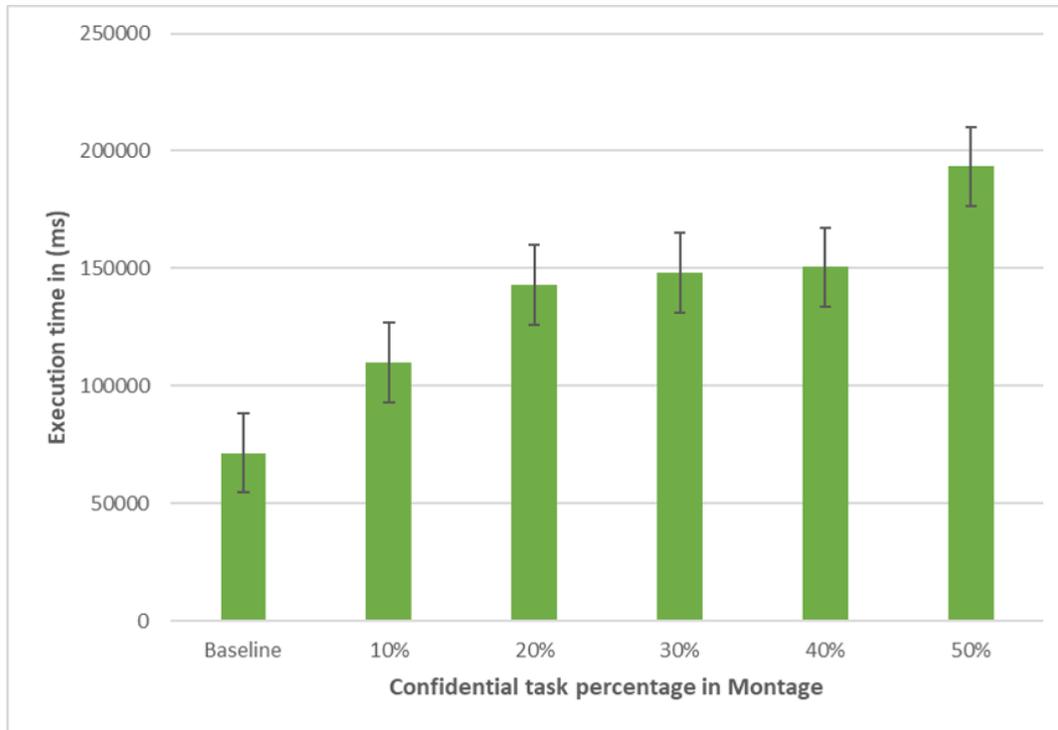
- The average performance overheads for the LIGO workflow shows a range between 1.81x in 10% and increases to 2.50x in 50% confidential task assignments.



LIGO Workflows

# Performance overhead of Montage workflow

- The average performance overheads for the Montage workflow shows a range between 1.34x in 10% and increases to 2.50x in 50% confidential task assignments.



Montage Workflows

# Conclusions and future work

- ▶ This paper presents an efficient, cost-effective deadline constrained algorithm SEED for scientific workflows in IaaS clouds.
- ▶ To the best of our knowledge, this is one of the first scheduling algorithms that can handle confidential workflow tasks.
- ▶ The primary intention is to schedule confidential tasks and achieve a low-cost scheduling algorithm within the given deadline while the complexity remains feasible for running large-scale workflows.
- ▶ For future works, we intend to protect the security of big data workflows in the heterogeneous cloud environments, including Intel SGX, AMD SEV [5], and GPU accelerators.

# SEED Demo

The screenshot displays the Eclipse IDE interface. The Package Explorer on the left shows a project named 'SEED' with a source folder 'src' containing 'SEED.java'. The main editor window shows the code for 'WorkflowPlanner\_SEED.java'. The code includes a method 'canSatisfy' that checks if a path of tasks can be satisfied by a machine, and a method 'splitCriticalPathByConfidentialTask' that splits a critical path into confidential tasks. The console output shows the execution of the SEED application, displaying task details and the resulting path and cost for each machine.

```
334
335
336
337  /* Checking whether each of the EFT meets the LFT*/
338  private boolean canSatisfy(ArrayList<Integer> path, int vm) {
339
340      double newStartTime = EST[path.get(0)];
341      for (int i = 0; i < path.size(); i++) {
342          int task = path.get(i);
343          double executionTime = execTime.get(vm + "").get(task);
344          /* We also need to check whether the new EST is at least or greater than previous EST*/
345          if (EST[task] <= newStartTime && newStartTime + executionTime <= LFT[task]) {
346              newStartTime += executionTime;
347          } else {
348              return false;
349          }
350      }
351      return true;
352  }
353
354  private ArrayList<ArrayList<Integer>> splitCriticalPathByConfidentialTask(ArrayList<Integer> criticalPath) {
355      ArrayList<ArrayList<Integer>> paths = new ArrayList<ArrayList<Integer>>();
356      if (confidentialTasks.isEmpty()) {
357          paths.add(criticalPath);
358      }
359      return paths;
360  }
```

```
<terminated> SEED [Java Application] /Library/Java/JavaVirtualMachines/jdk-13.0.2.jdk/Contents/Home/bin/java (Oct 9, 2020, 3:20:52 PM)
2
3
1
6
5
4
7
8
9
Task 1 0.0 8.0 9.0
Task 2 0.0 12.0 14.0
Task 3 0.0 9.0 12.0
Task 4 8.0 18.0 19.0
Task 5 14.0 22.0 24.0
Task 6 12.0 20.0 22.0
Task 7 18.0 29.0 30.0
Task 8 22.0 28.0 30.0
Task 9 20.0 28.0 30.0
Path :
2 6 9 Assigned machine : 1 Cost : 6.0
Path :
3 Assigned machine : 2 Cost : 1.0
Path :
1 4 7 Assigned machine : 2 Cost : 3.0
Path :
5 8 Assigned machine : 1 Cost : 4.0
Total cost : 14.0
```

# SecDATAVIEW Introduction

- ▶ SecDATAVIEW [4] system is our advance design of DATAVIEW that protects the confidentiality and privacy of big data workflows in the heterogeneous cloud environments.
- ▶ SecDATAVIEW can leverage hardware-based security features such as Intel SGX [2] and AMD SEV [5] and guarantee the secure execution of confidential workflow in the untrusted cloud environments.
- ▶ Access DATAVIEW and SecDATAVIEW:
  - <https://github.com/shiyonglu/DATAVIEW>
  - <https://github.com/shiyonglu/SecDATAVIEW>
  - DATAVIEW YouTube Channel:
  - <https://www.youtube.com/channel/UCrIEBUmju-NMKFMIFbsKBYw>

# DATAVIEW TEAM





# Thank You.

# Questions?

# References:

- ▶ [1] McKeen *et al.*, “Innovative instructions and software model for isolated execution.,” in *HASP@ISCA*, 2013, p. 10.
- ▶ [2] A. Kashlev, S. Lu, and A. Mohan, “Big Data Workflows: a Reference Architecture and the DATAVIEW System,” *Serv. Trans. Big Data*, vol. 4, no. 1, pp. 1–19, 2017.
- ▶ [3] Ahmed *et al.*, “Diagnosis Recommendation using Machine Learning Scientific Workflows,” in *Big Data Congress, 2018 IEEE International Conference on*, 2018.
- ▶ [4] Mofrad, S., Ahmed, I., Lu, S., Yang, P., Cui, H., & Zhang, F. (2019, December). SecDATAVIEW: a secure big data workflow management system for heterogeneous computing environments. In *Proceedings of the 35th Annual Computer Security Applications Conference* (pp. 390–403).
- ▶ [5] Kaplan, D., Powell, J., & Woller, T. (2016). AMD memory encryption. *White paper*.