IEEE/ACM International Conference on
Big Data Computing, Applications and Technologies

Nantes, France
02 Dec, 2025

An Empirical Study on

Estimating Energy Consumption From Software Traces

Authors: Phuong Nguyen[1,2] *(presenter)*, Adel Noureddine[3,4] and Congduc Pham[1,2]

(1) Université de Pau et des Pays de l'Adour
(2) LIUPPA
(3) Université Paris Nanterre
(4) LIP6

This research was supported by CCLO (France) and the RESILINK project (PRIMA S2 2021, EU-funded).

RESEARCH CONTEXT

RESEARCH CONTEXT

- ICT – A Double-Edged Sword

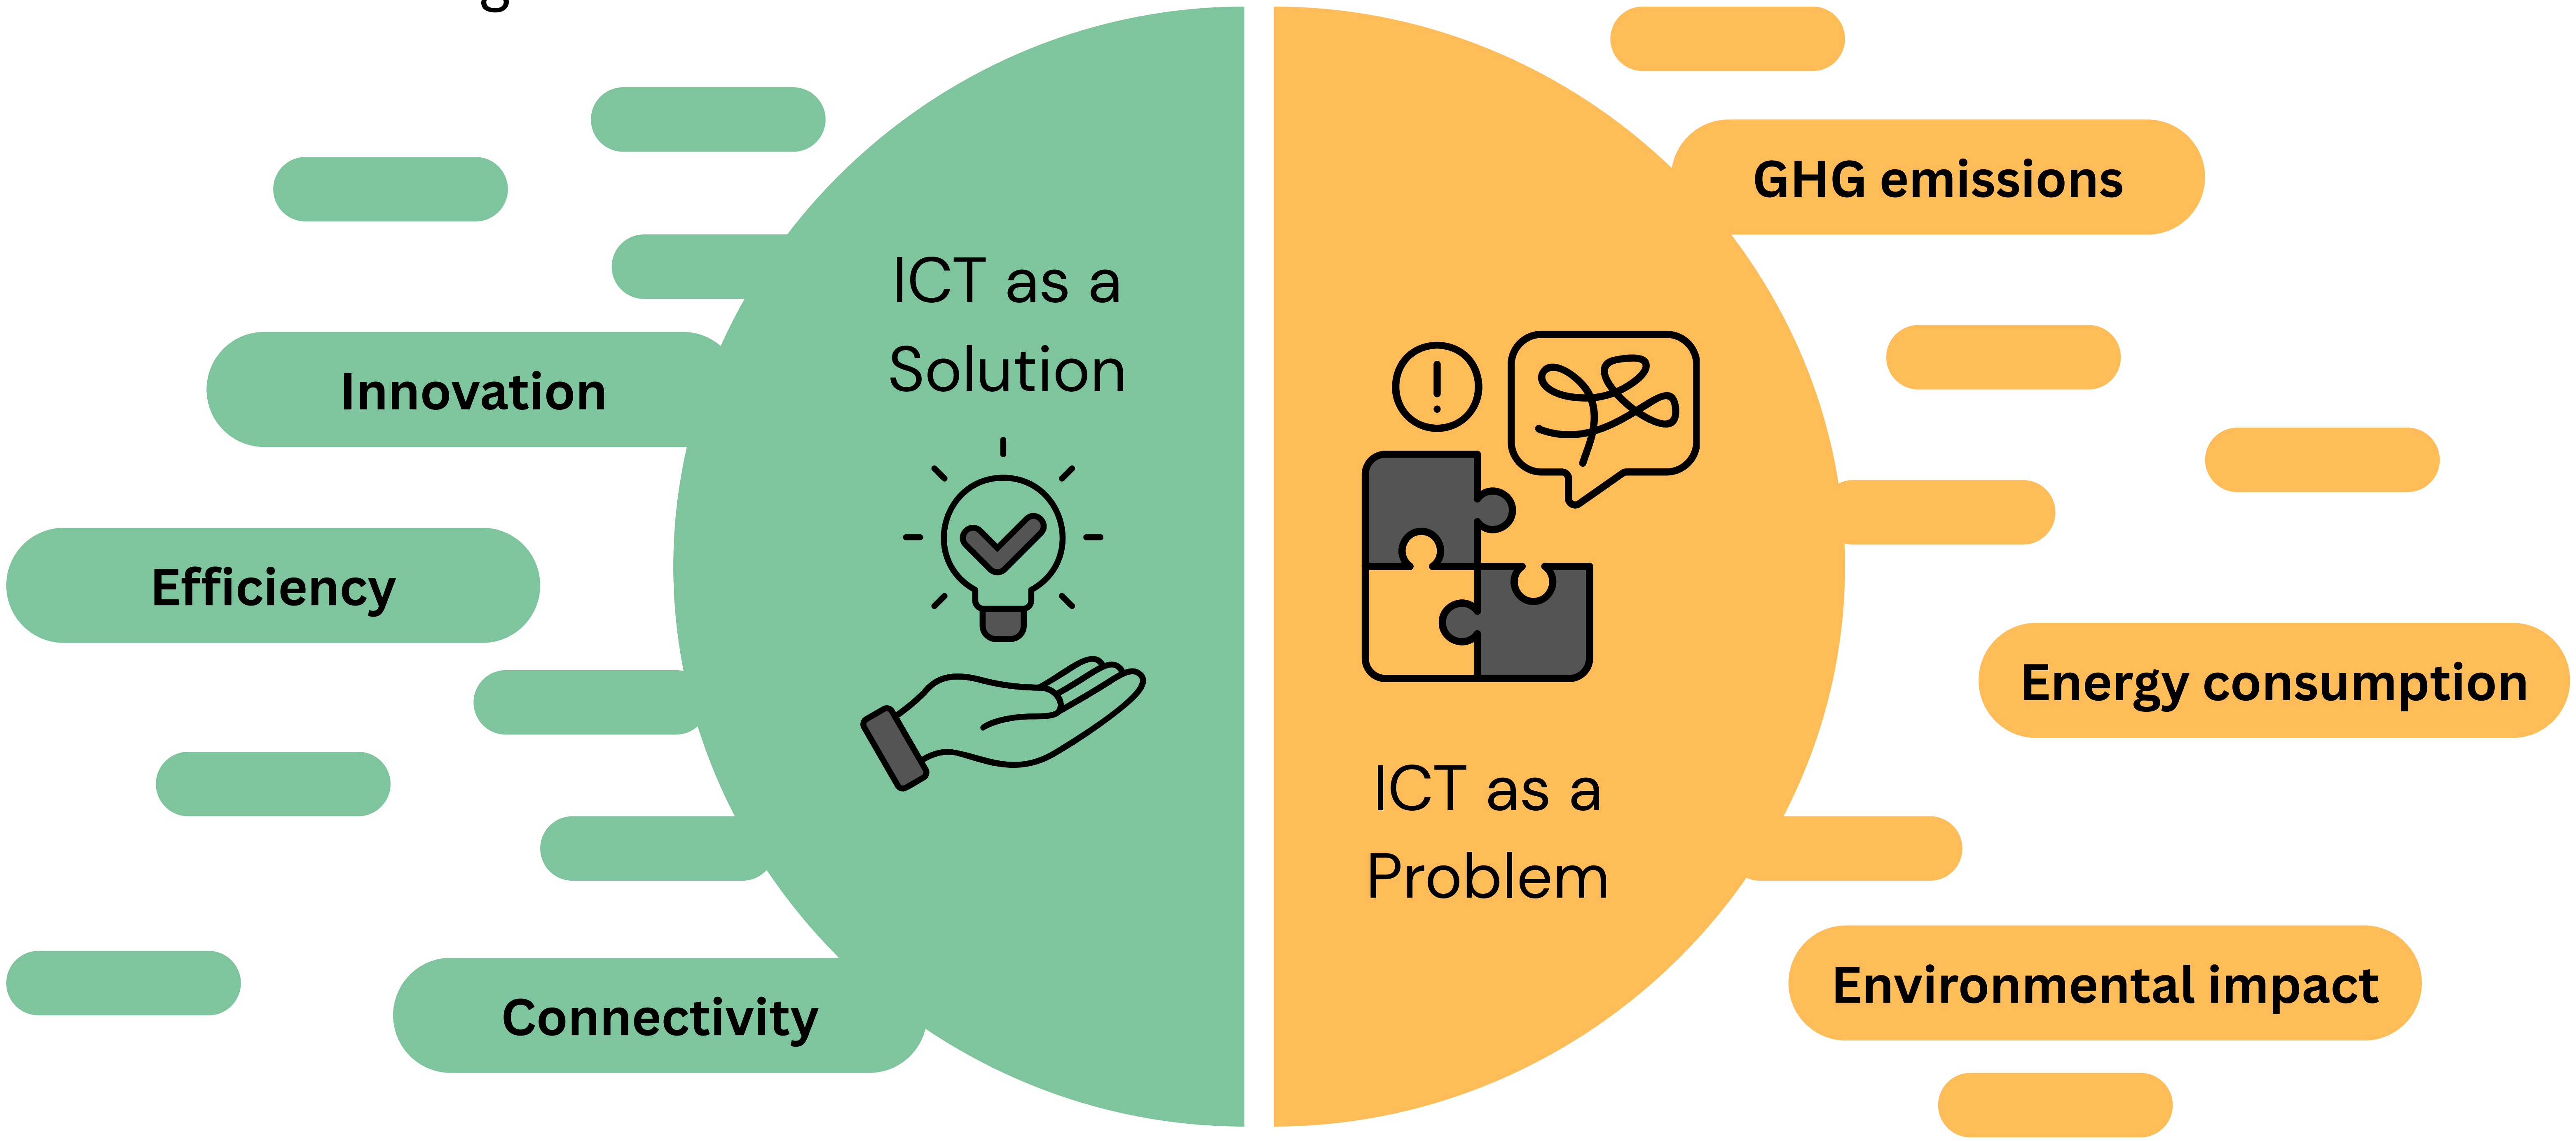

RESEARCH CONTEXT

- The Legacy Software Challenge

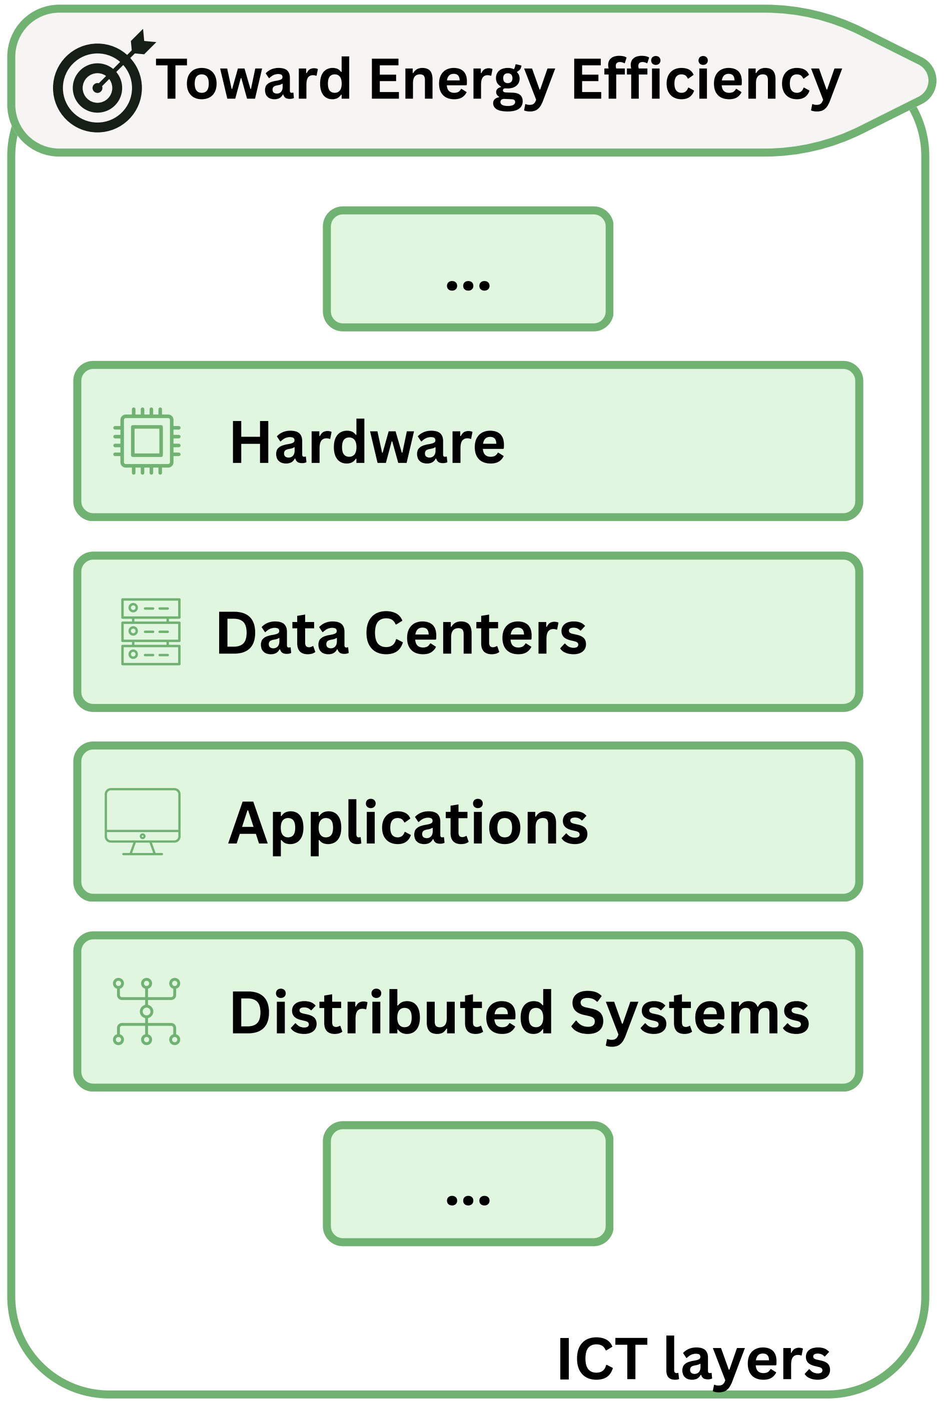

RESEARCH CONTEXT

- The Legacy Software Challenge

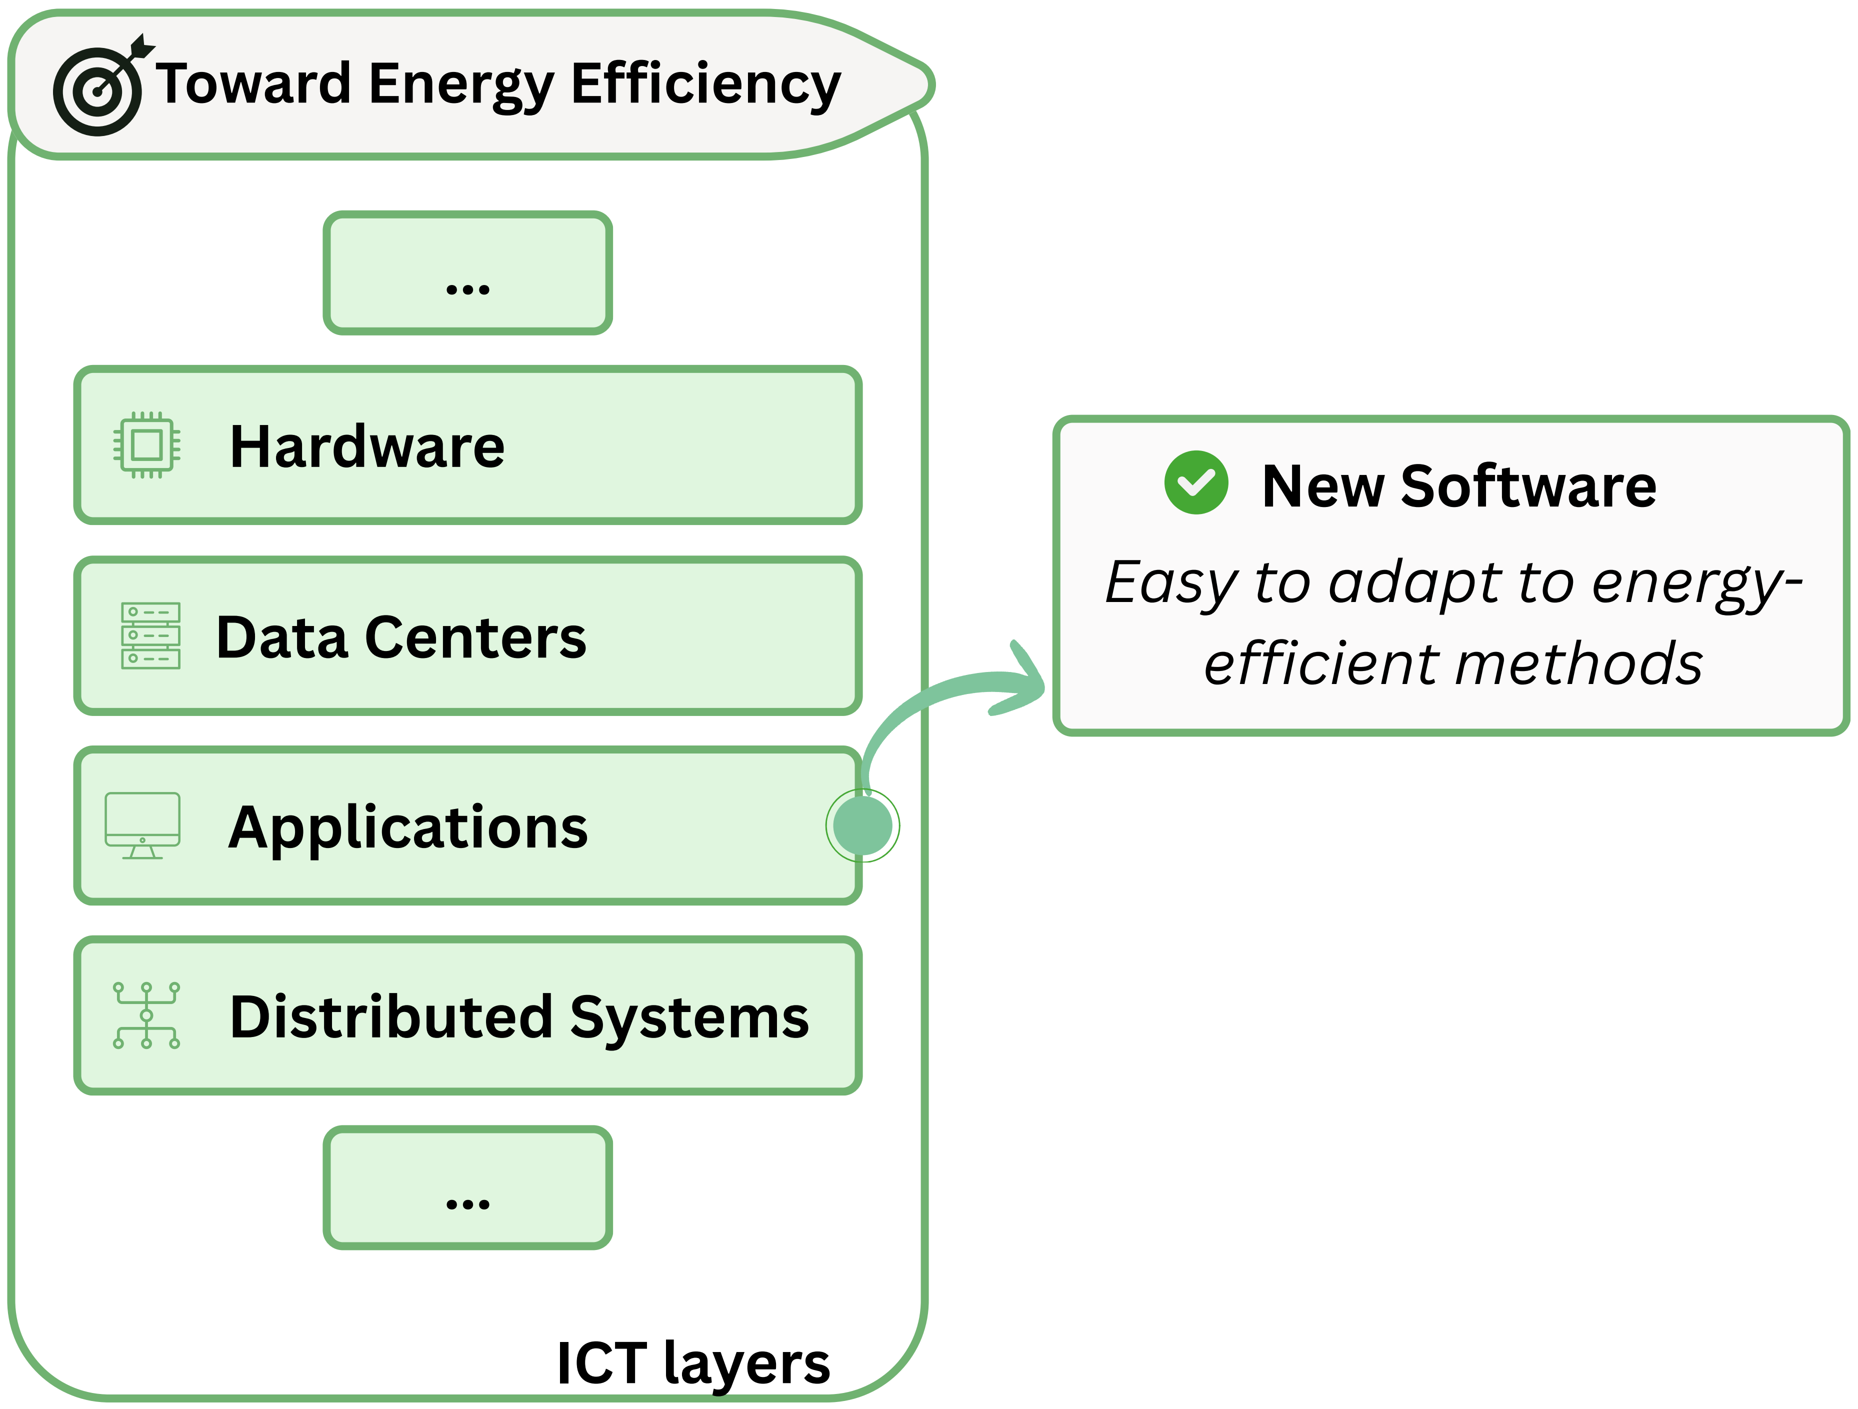

RESEARCH CONTEXT

- The Legacy Software Challenge

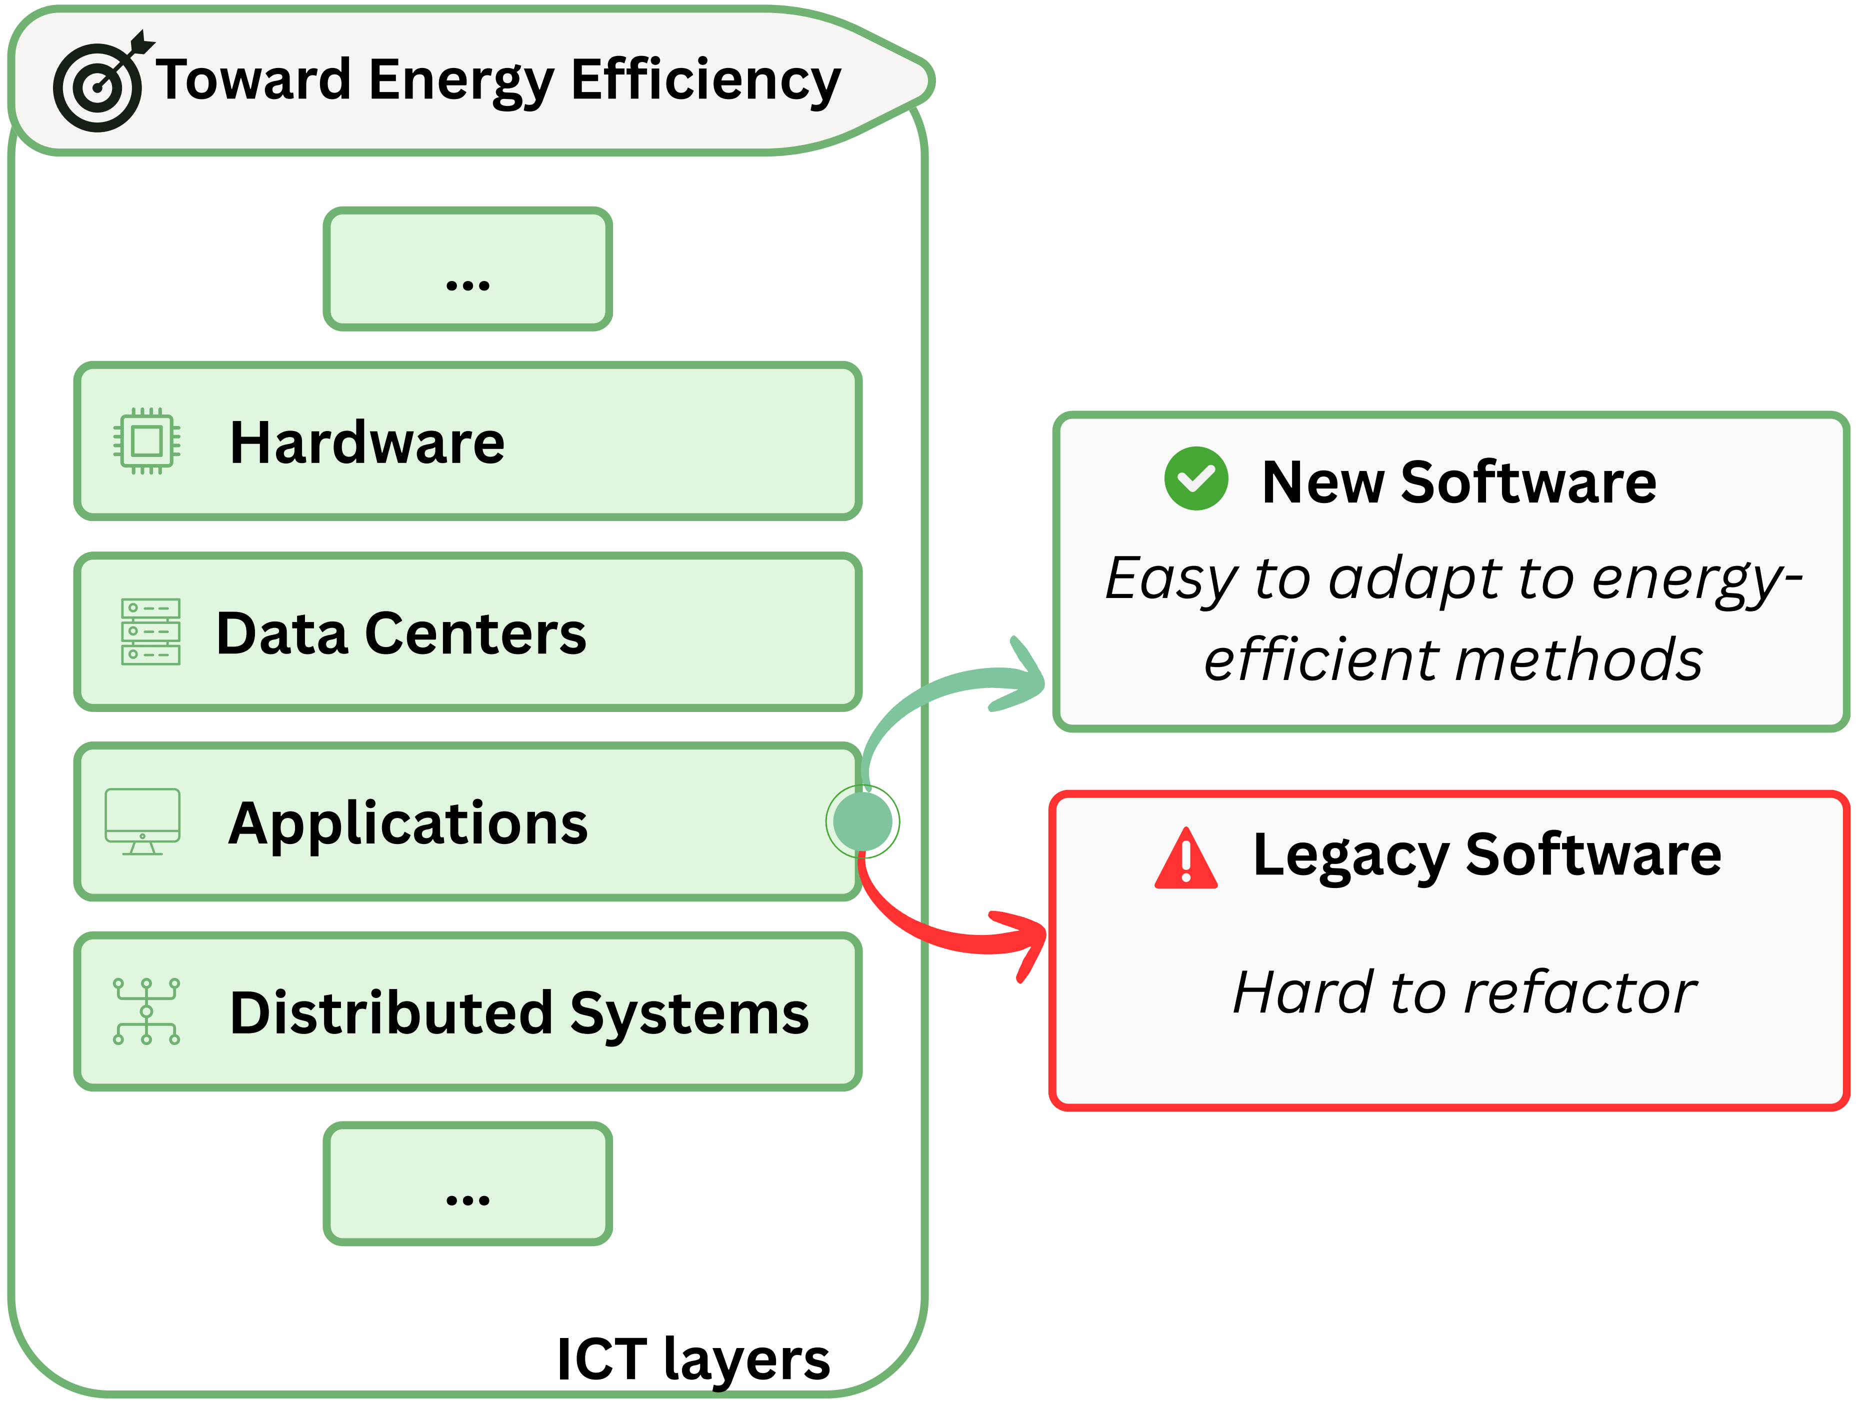

RESEARCH CONTEXT

- The Legacy Software Challenge

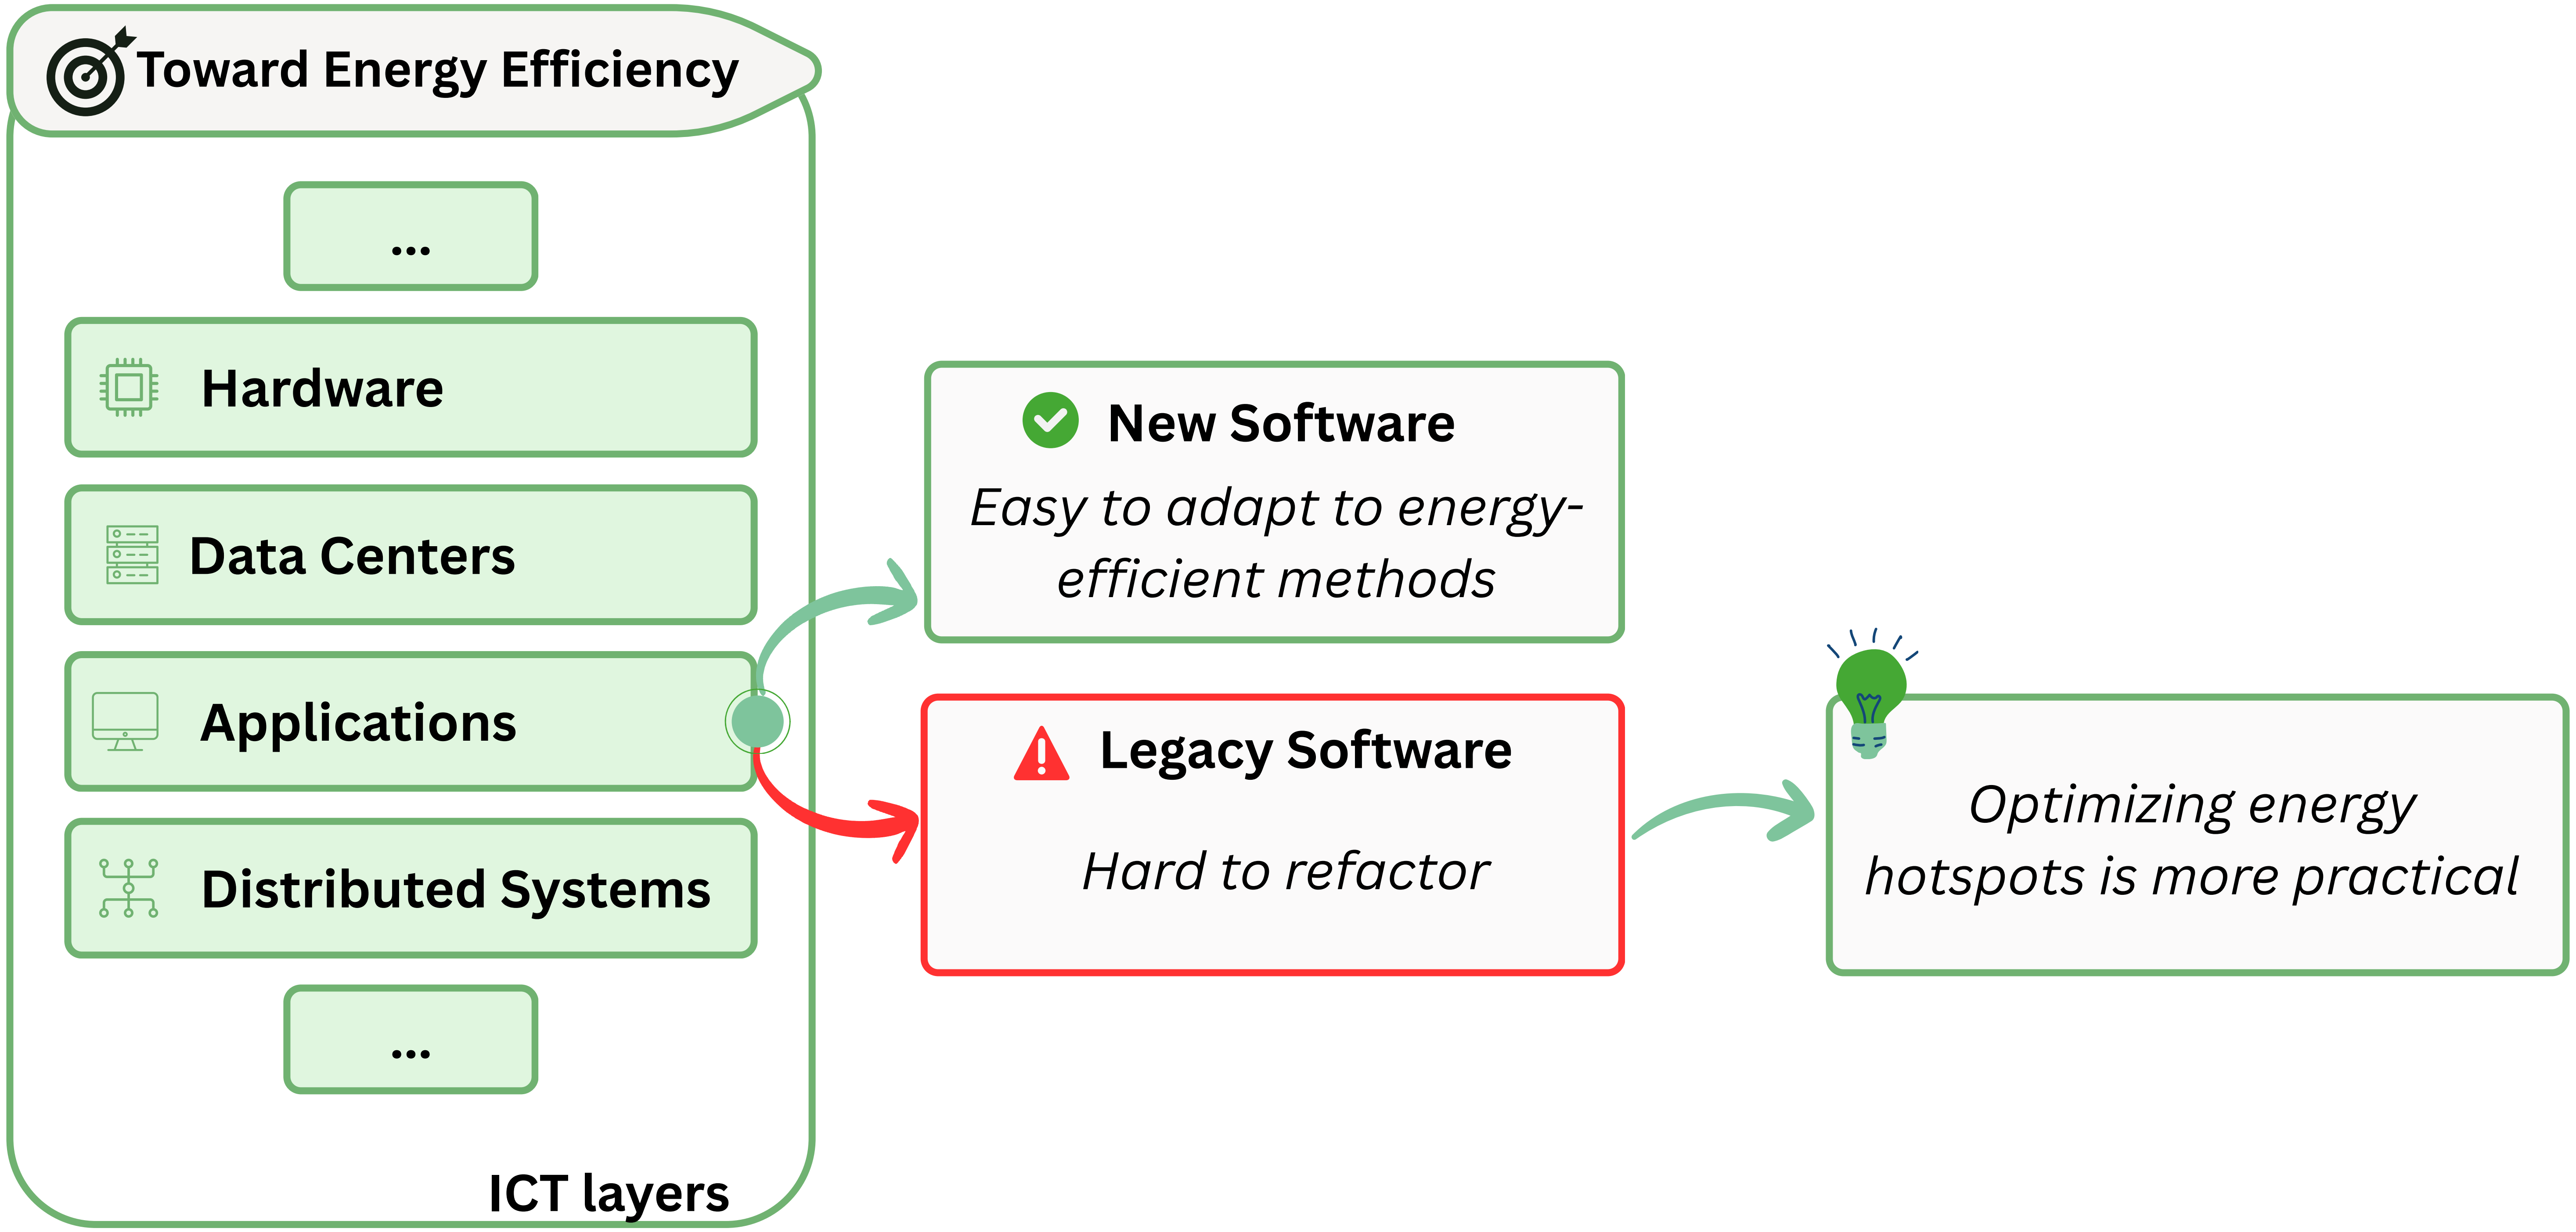

RESEARCH CONTEXT

- Optimizing Energy Hotspots

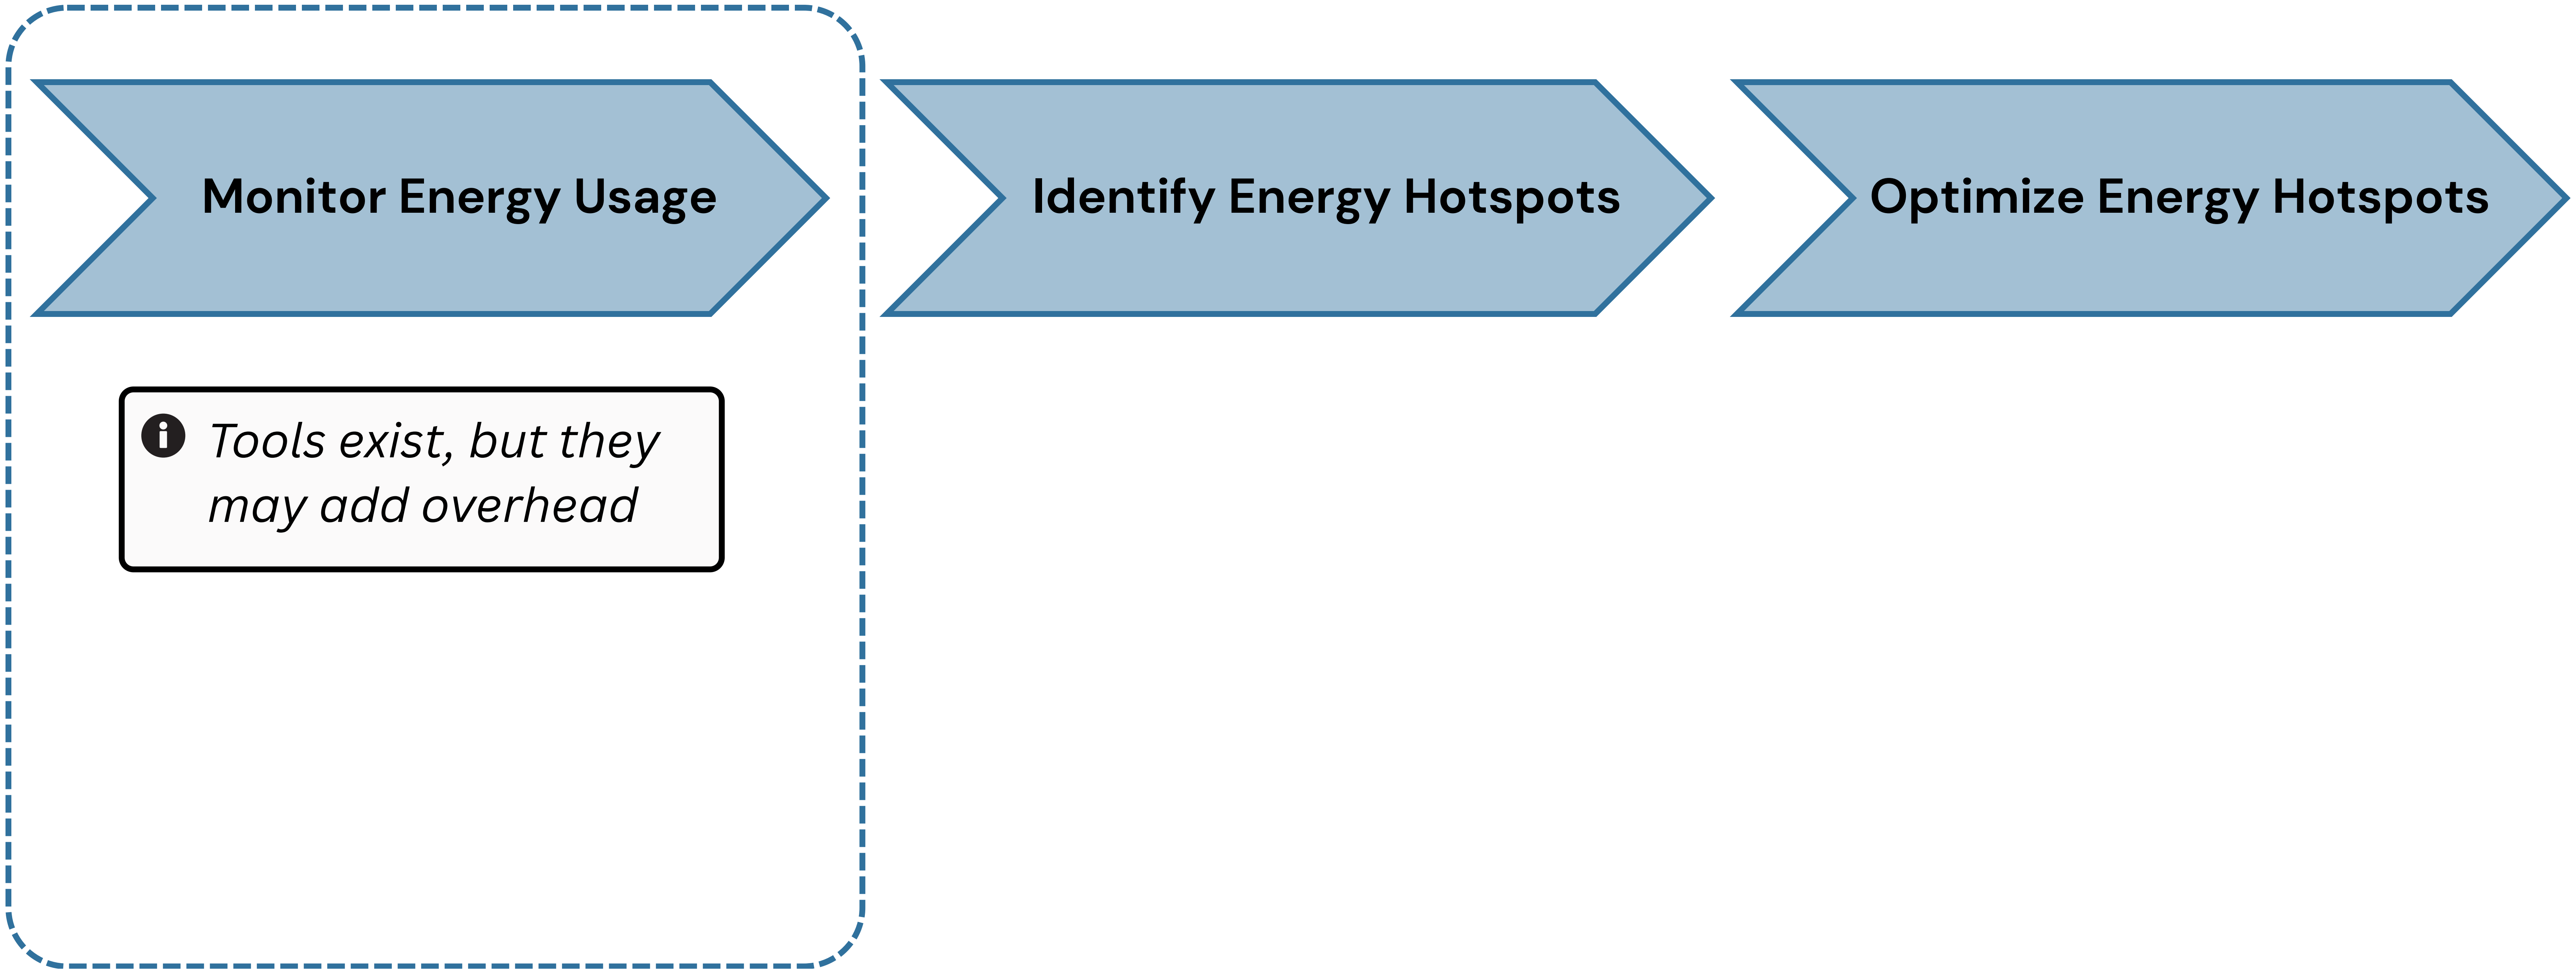

RESEARCH CONTEXT

- Optimizing Energy Hotspots

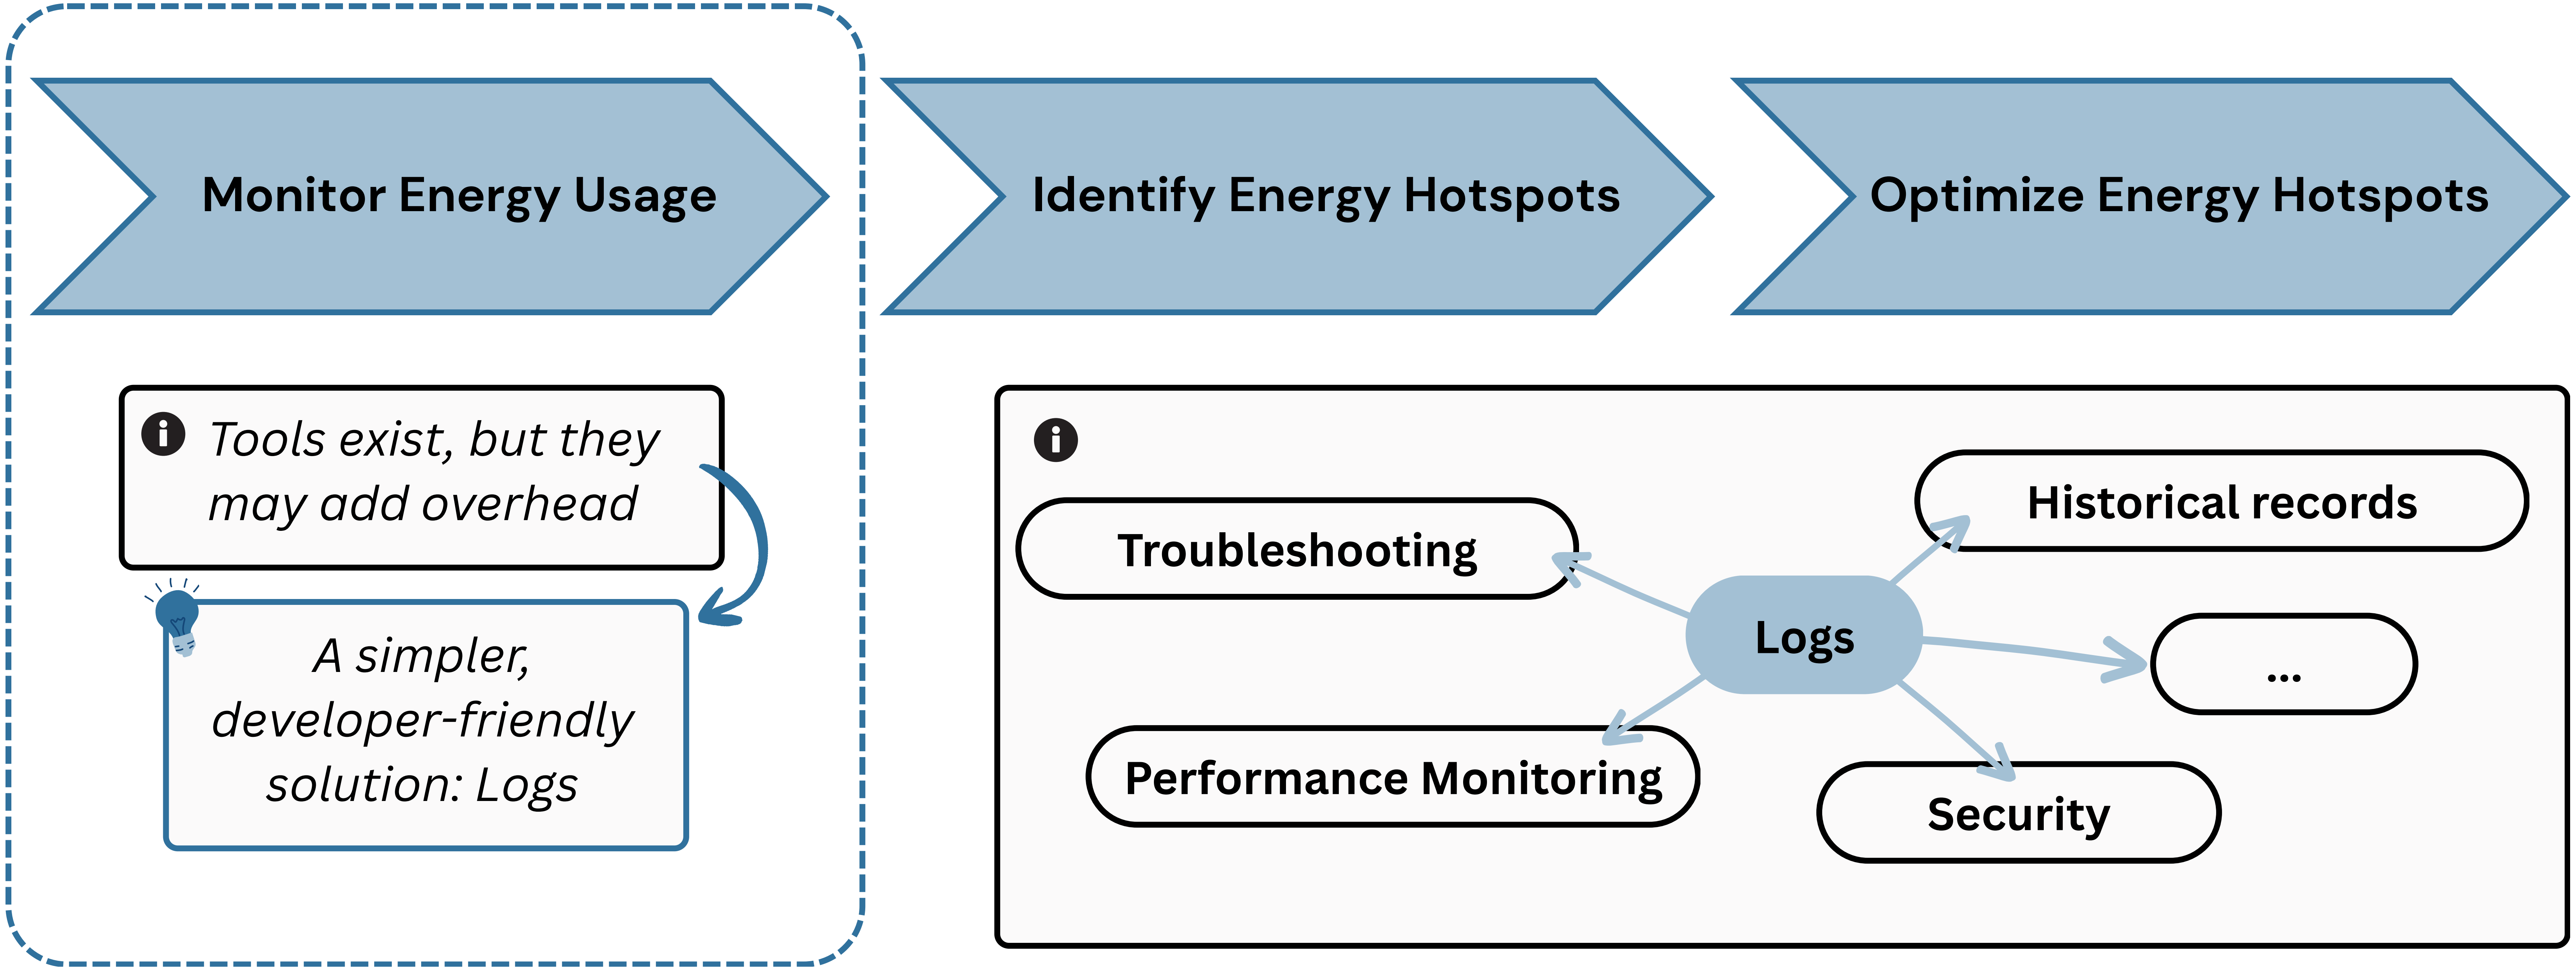

RESEARCH CONTEXT

- Optimizing Energy Hotspots

Monitor Energy Usage → **Identify Energy Hotspots** → **Optimize Energy Hotspots**

Tools exist, but they may add overhead

A simpler, developer-friendly solution: Logs

Research Question:

Can software energy consumption be estimated using software traces and logs?

METHODOLOGY OVERVIEW

METHODOLOGY OVERVIEW

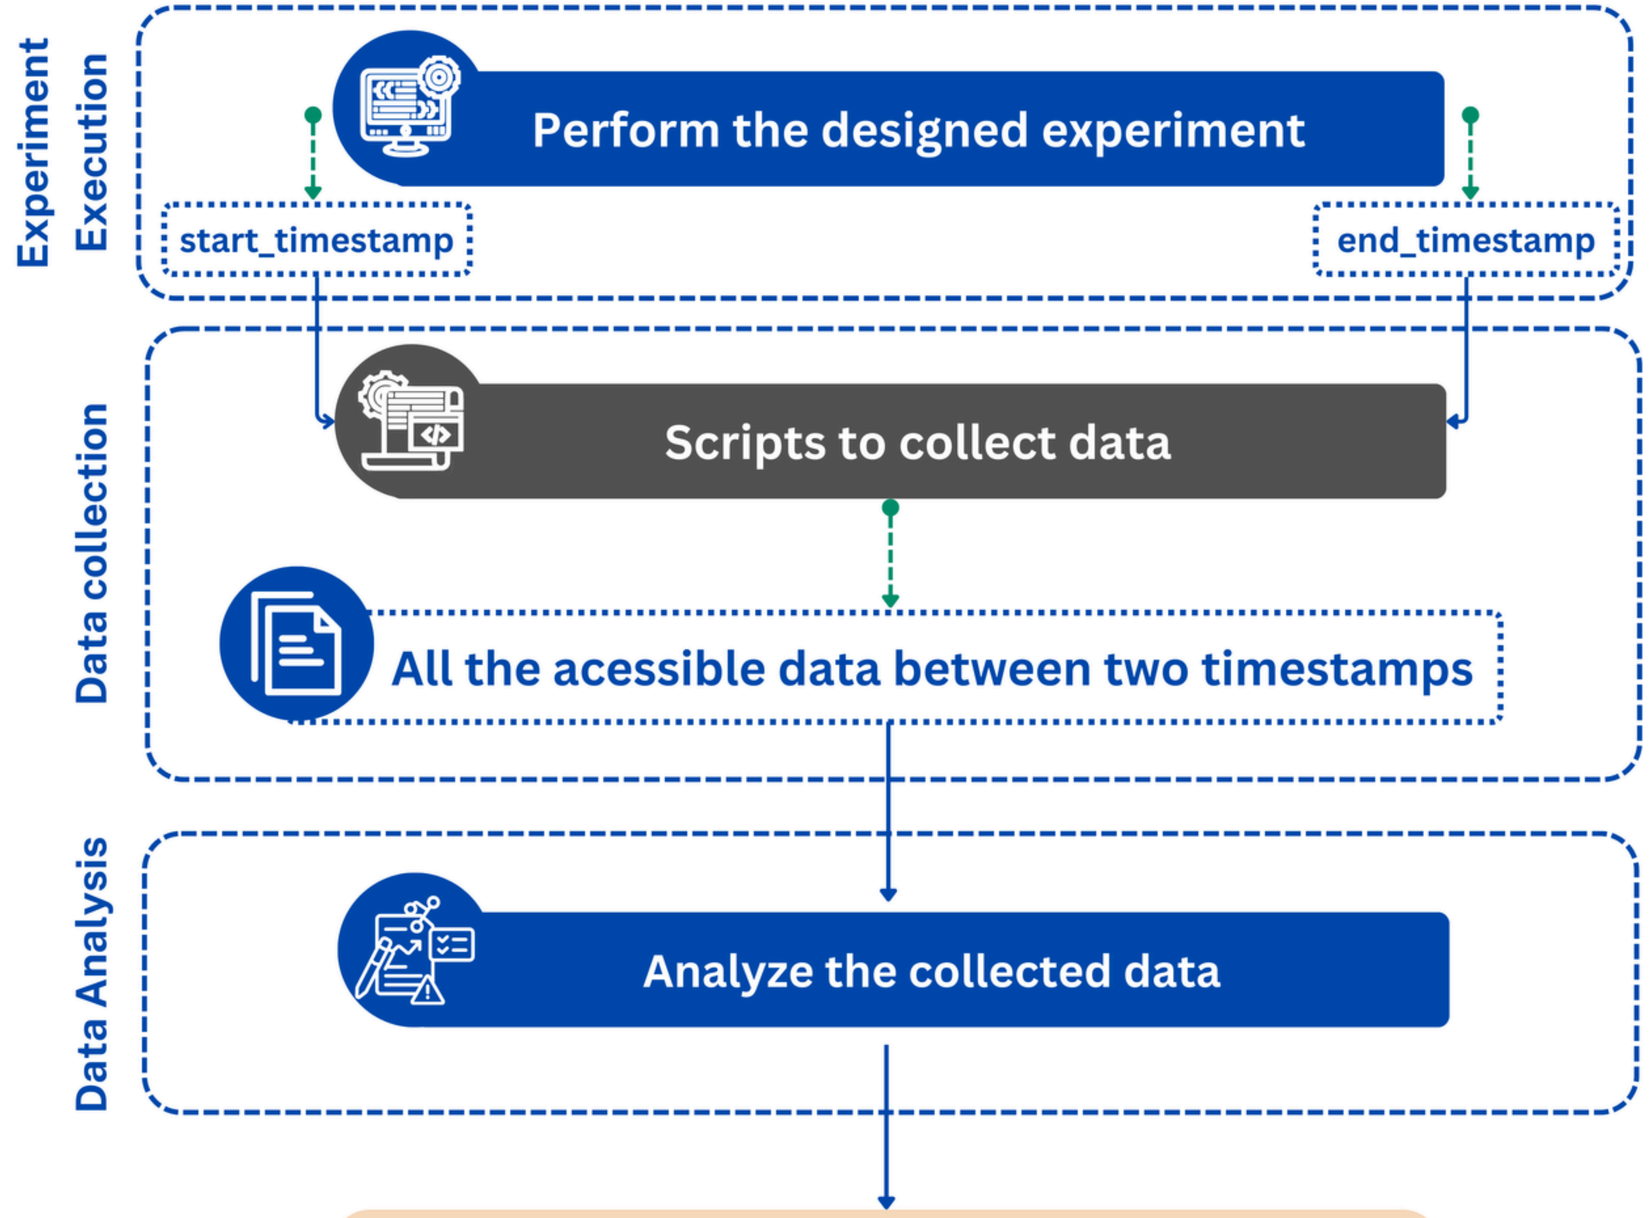

RQ: Can *software energy consumption* be estimated using *software traces* and *logs*?

METHODOLOGY OVERVIEW

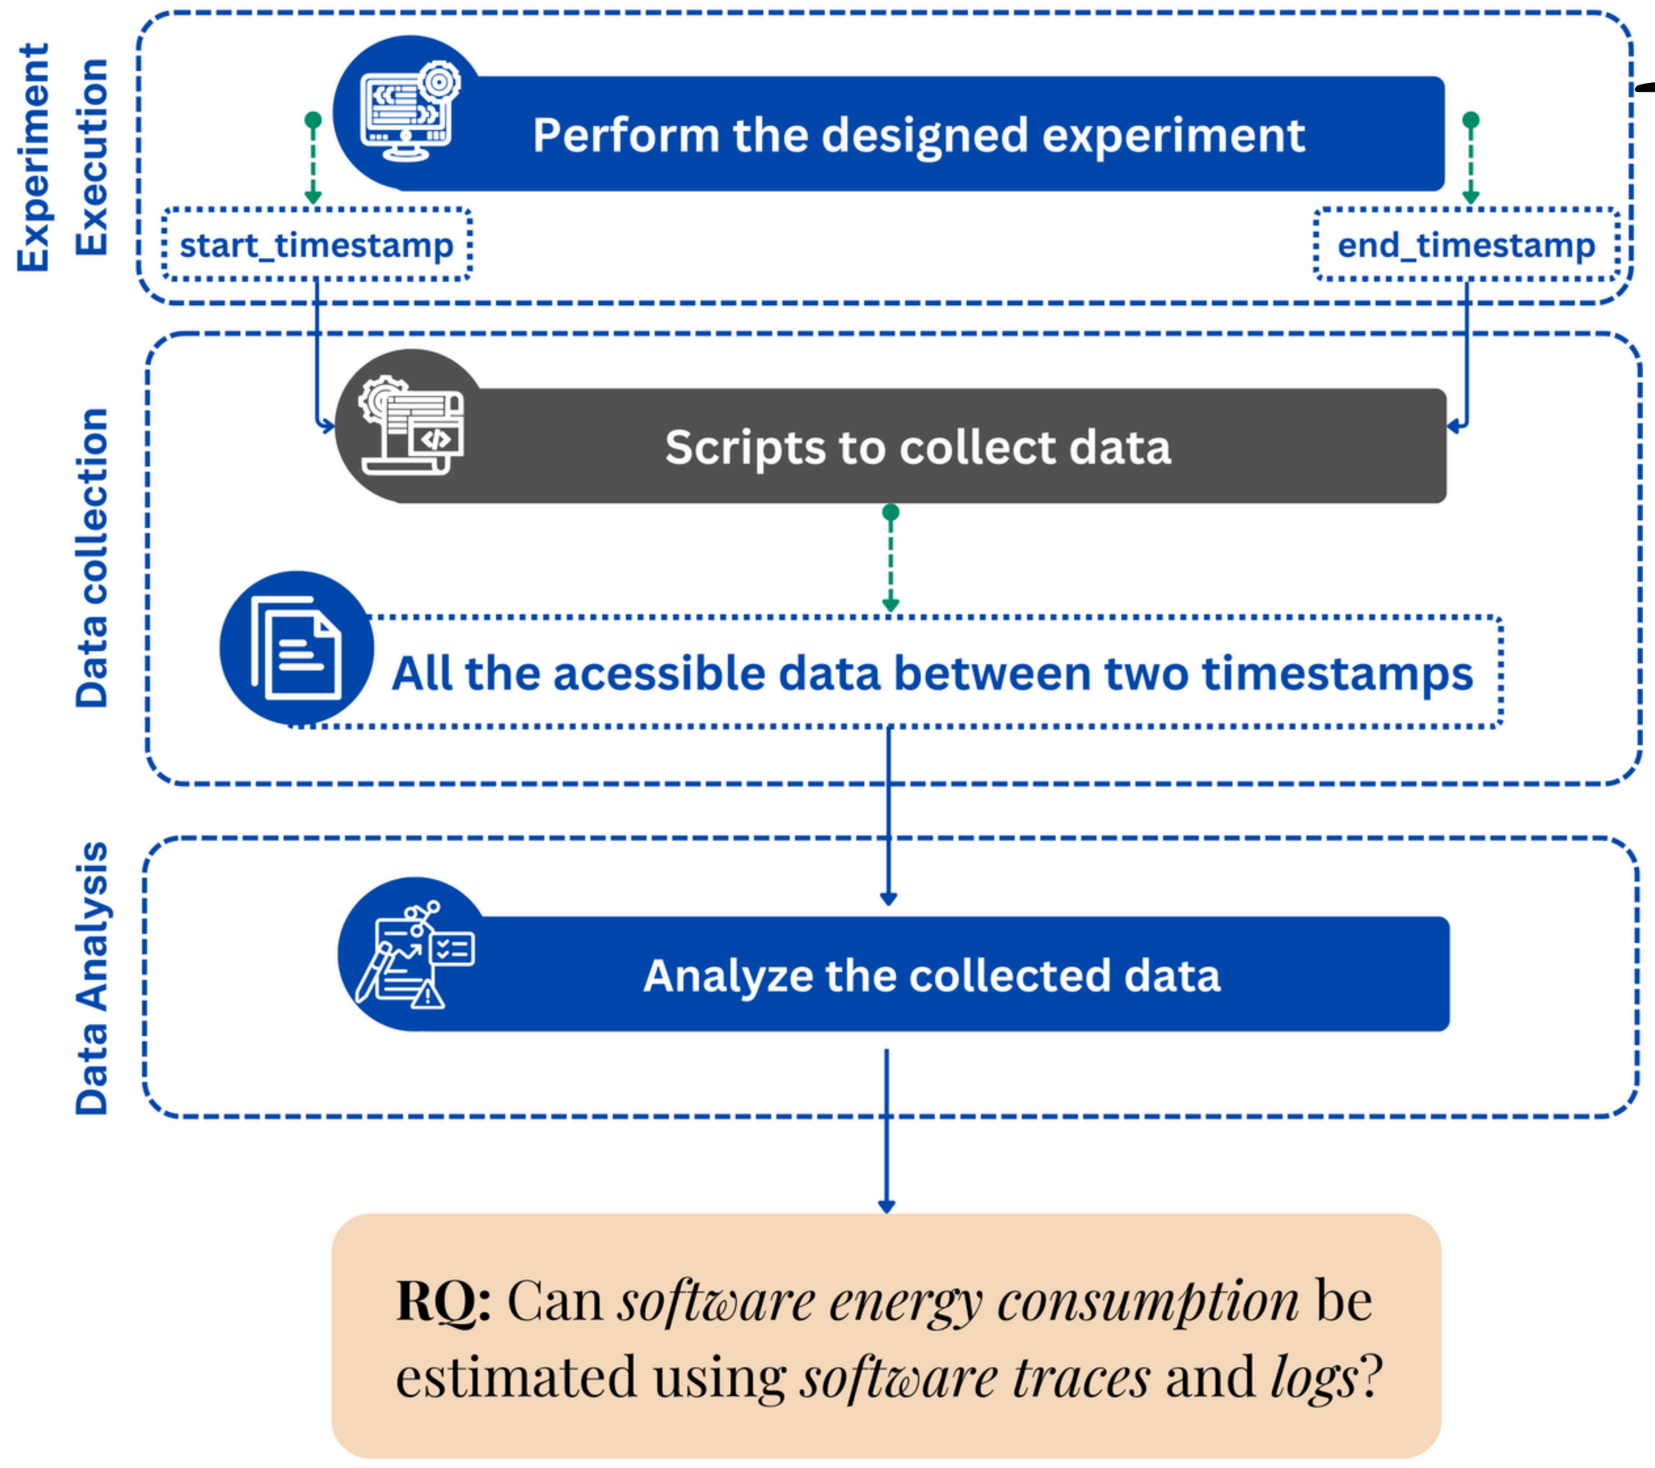

- Two Types of Workloads

Synthetic Workloads

CPU-intensive *memory-intensive* *combined tasks*

Realistic Workloads

DayTrader *Hotel Reservation* *Maven*
Custom REST App

- Run each workload **multiple times** for stability & reliability

METHODOLOGY OVERVIEW

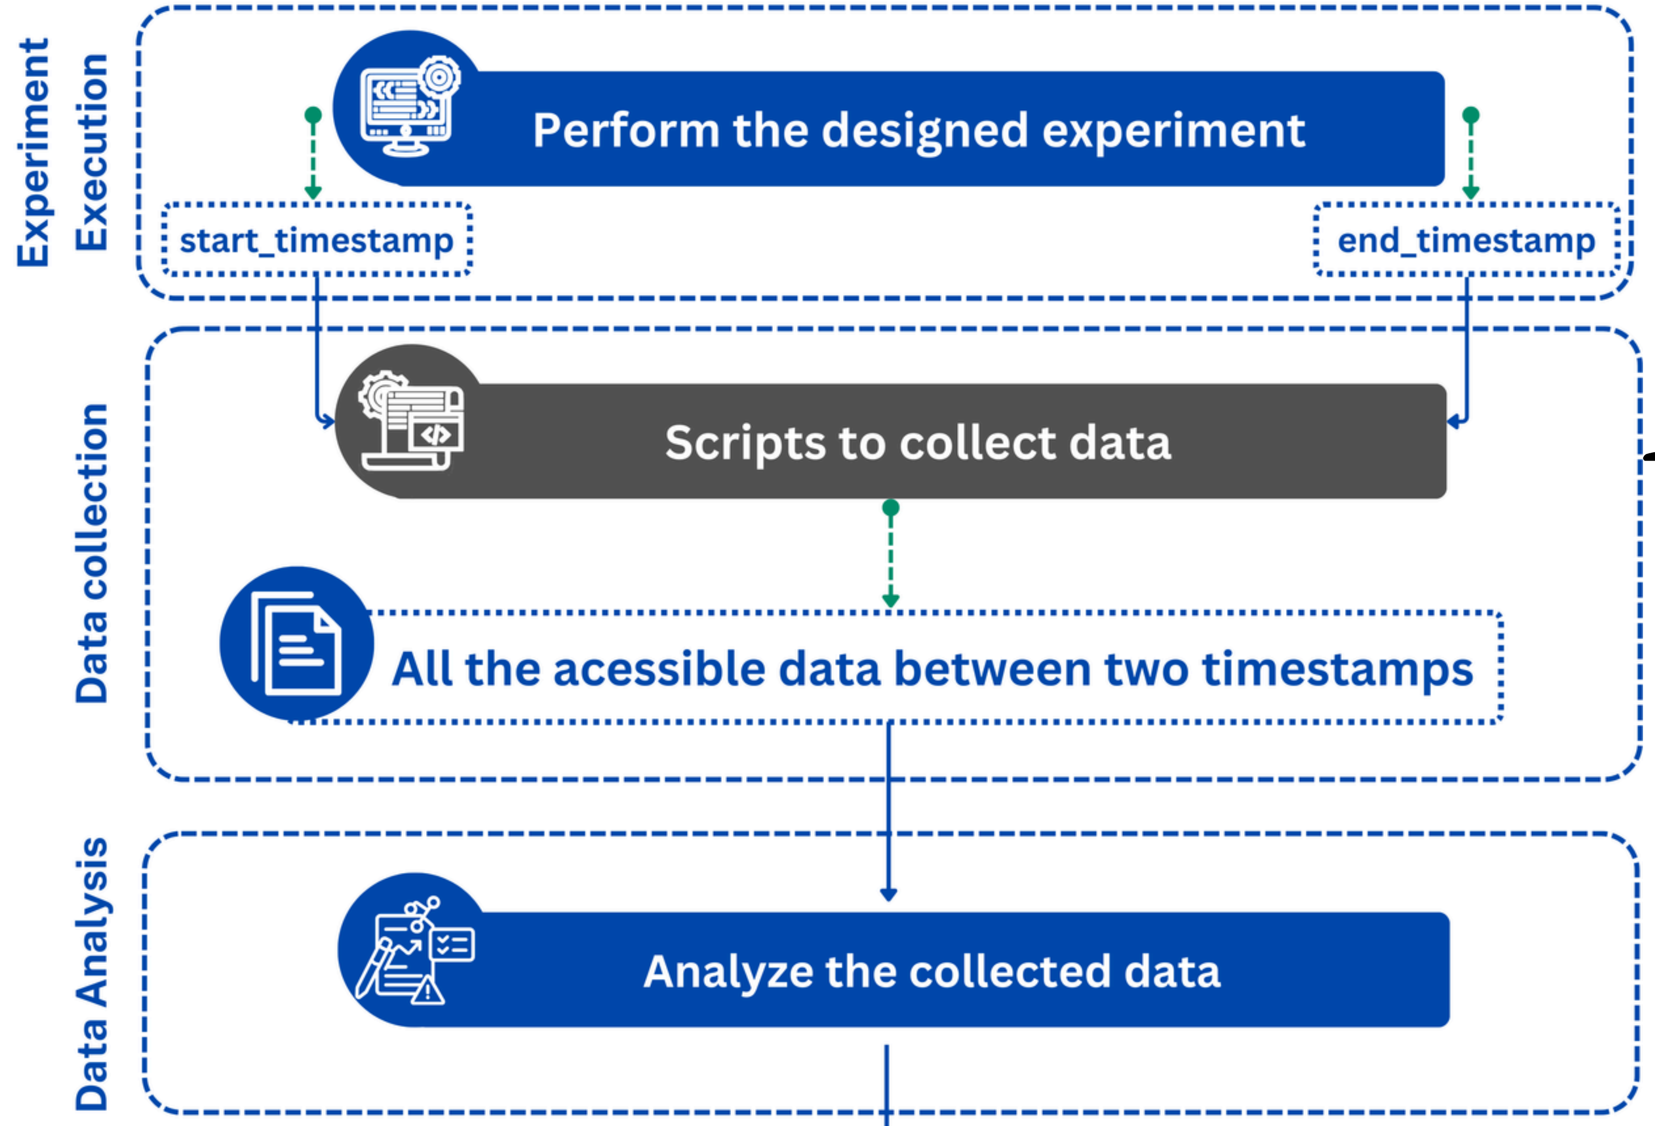

RQ: Can *software energy consumption* be estimated using *software traces* and *logs*?

- A **variety of metrics** from multiple sources.
  For example:

  Hardware metrics | Database Logs | Maven logs | Docker Logs | Network metrics

  ** The type and scope of these metrics varied depending on the workload*

- Energy Measurement: **Intel RAPL**

METHODOLOGY OVERVIEW

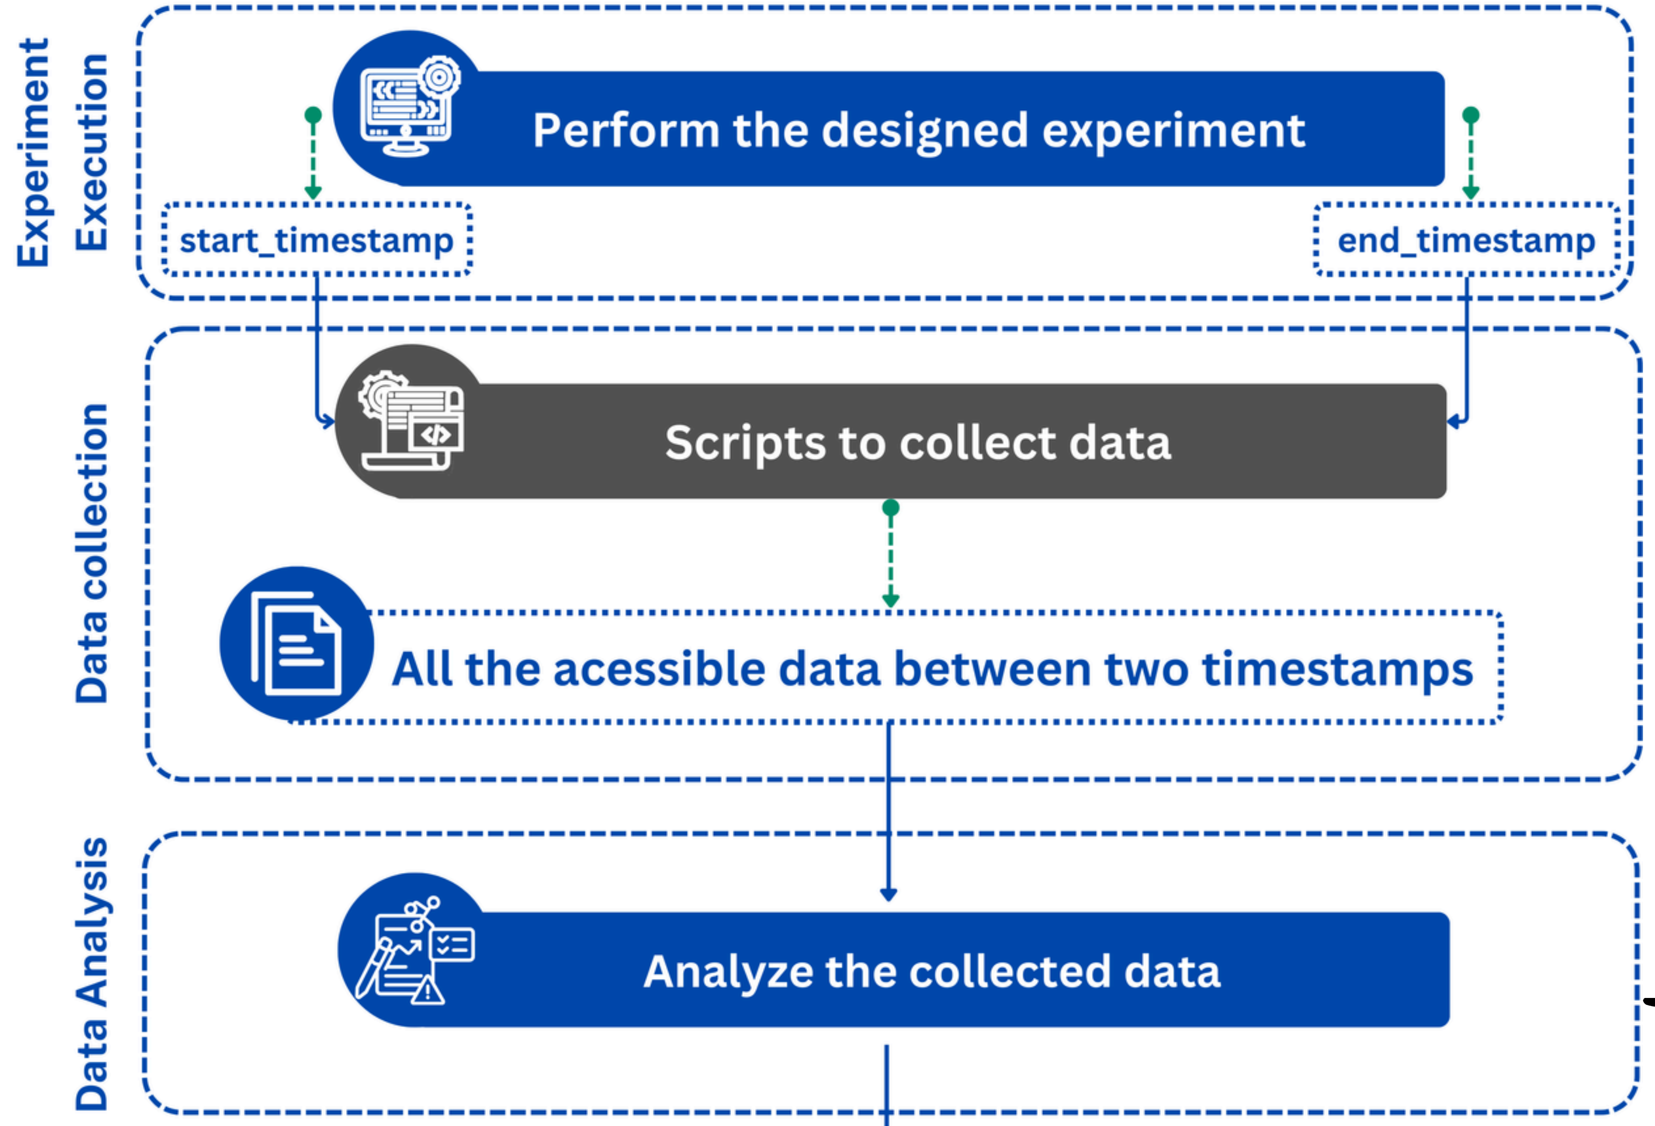

RQ: Can *software energy consumption* be estimated using *software traces* and *logs*?

- Compute **Pearson correlations** (with p-values) between **energy usage** and **each metric**.

** Compute correlations per run, then average across runs*

RESULTS & INSIGHTS

RESULTS & INSIGHTS – Synthetic Workloads

** Only statistically significant correlations are reported: $|r| > 0.3$, $p < 0.05$*

Results

- **CPU utilization**: r = 0.452 - 0.739, $p < 0.001$
- **Disk writes**: r = 0.462 - 0.534, $p < 0.001$ in CPU-intensive workload.
- **Memory utilization** & **disk reads**: not significant

Table 1: Pearson correlations between collected metrics and energy consumption for synthetic workloads.

	100ms interval			1s interval		
	(i)	(ii)	(iii)	(i)	(ii)	(iii)
Collected metrics	r	r	r	r	r	r
CPU Utilization	0.489*	0.582*	0.452*	0.547*	0.739*	0.505*
#Reads on disk	0.249*	0.086	0.026	0.247*	0.060	-0.068
#Writes on disk	0.462*	0.132*	0.016	0.534*	0.129	-0.076
Memory Utilization	-0.092*	0.144*	-0.007	-0.122	0.153*	-0.085

(i) CPU-intensive; (ii) Memory-intensive; (iii) Combined tasks
r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS – Synthetic Workloads

** Only statistically significant correlations are reported: $|r| > 0.3$, $p < 0.05$*

Results

- **CPU utilization**: r = 0.452 - 0.739, $p < 0.001$
- **Disk writes**: r = 0.462 - 0.534, $p < 0.001$ in CPU-intensive workload.
- **Memory utilization** & **disk reads**: not significant

Insights

- **CPU utilization** - **strong positive predictor** of energy consumption
- **Disk writes** - *matter only in CPU-heavy tasks*
- **Memory utilization** - *has minimal impact*

Table 1: Pearson correlations between collected metrics and energy consumption for synthetic workloads.

	100ms interval			1s interval		
	(i)	(ii)	(iii)	(i)	(ii)	(iii)
Collected metrics	r	r	r	r	r	r
CPU Utilization	0.489*	0.582*	0.452*	0.547*	0.739*	0.505*
#Reads on disk	0.249*	0.086	0.026	0.247*	0.060	-0.068
#Writes on disk	0.462*	0.132*	0.016	0.534*	0.129	-0.076
Memory Utilization	-0.092*	0.144*	-0.007	-0.122	0.153*	-0.085

(i) CPU-intensive; (ii) Memory-intensive; (iii) Combined tasks
r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS - DayTrader

** Only statistically significant correlations are reported: $|r| > 0.3$, $p < 0.05$*

Results – DayTrader 7

- **# Network packets**: $r = -0.335$, $p = 0.001$
- **Total packet length**: $r = -0.321$, $p = 0.001$
- **# HTTP requests**: $r = -0.314$, $p = 0.001$

Table 2: Pearson correlations between collected metrics and energy consumption for DayTrader.

	DayTrader 7	DayTrader 10
Collected metrics	r	r
#HTTP Requests	-0.314**	-0.25*
#Req. by Method GET	-0.313**	-0.25*
#Req. by Response 200	-0.316**	-0.251*
#Network Packets	-0.335**	-0.242*
Total Length of Packets	-0.321**	-0.241*
AS Memory Utilization	0.243	0.432*

r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS - DayTrader

* *Only statistically significant correlations are reported: $|r| > 0.3$, $p < 0.05$*

Results – DayTrader 7

- **# Network packets**: $r = -0.335$, $p = 0.001$
- **Total packet length**: $r = -0.321$, $p = 0.001$
- **# HTTP requests**: $r = -0.314$, $p = 0.001$

Results – DayTrader 10

- **Application server memory utilization**: $r = 0.432$, $p = 0.012$

Table 2: Pearson correlations between collected metrics and energy consumption for DayTrader.

	DayTrader 7	DayTrader 10
Collected metrics	r	r
#HTTP Requests	-0.314**	-0.25*
#Req. by Method GET	-0.313**	-0.25*
#Req. by Response 200	-0.316**	-0.251*
#Network Packets	-0.335**	-0.242*
Total Length of Packets	-0.321**	-0.241*
AS Memory Utilization	0.243	0.432*

r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS - DayTrader

** Only statistically significant correlations are reported: |r| > 0.3, p < 0.05*

Results – DayTrader 7

- **# Network packets**: $r = -0.335$, $p = 0.001$
- **Total packet length**: $r = -0.321$, $p = 0.001$
- **# HTTP requests**: $r = -0.314$, $p = 0.001$

Results – DayTrader 10

- **Application server memory utilization**: $r = 0.432$, $p = 0.012$

Insights

- **Workload variability** affects **energy predictability**

Table 2: Pearson correlations between collected metrics and energy consumption for DayTrader.

	DayTrader 7	DayTrader 10
Collected metrics	r	r
#HTTP Requests	-0.314**	-0.25*
#Req. by Method GET	-0.313**	-0.25*
#Req. by Response 200	-0.316**	-0.251*
#Network Packets	-0.335**	-0.242*
Total Length of Packets	-0.321**	-0.241*
AS Memory Utilization	0.243	0.432*

r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS – Hotel Reservation

** Only statistically significant correlations are reported: $|r| > 0.3$, $p < 0.05$*

Results

- **# Docker logs**: $r = 0.592$, $p < 0.001$
- **# Slow query logs**: $r = 0.716$, $p < 0.001$
- **# Logs from *reservation***: $r = 0.710$, $p < 0.001$
- **# Slow query logs from *reservation***: $r = 0.716$, $p < 0.001$

Table 3: Pearson correlations between collected metrics and energy consumption for HotelReservation.

Collected metrics	r
#Docker logs	0.592***
#Logs from reservation	0.710***
#Slow query logs	0.716***
#Slow query logs from reservation	0.716***

r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS – Hotel Reservation

** Only statistically significant correlations are reported: $|r| > 0.3$, $p < 0.05$*

Results

- **# Docker logs**: $r = 0.592$, $p < 0.001$
- **# Slow query logs**: $r = 0.716$, $p < 0.001$
- **# Logs from *reservation***: $r = 0.710$, $p < 0.001$
- **# Slow query logs from *reservation***: $r = 0.716$, $p < 0.001$

Insights

- **Structured logs** provide **more reliable** energy insight
- **The reservation container** acts as **the system bottleneck**

Table 3: Pearson correlations between collected metrics and energy consumption for HotelReservation.

Collected metrics	r
#Docker logs	0.592^{***}
#Logs from reservation	0.710^{***}
#Slow query logs	0.716^{***}
#Slow query logs from reservation	0.716^{***}

r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS – Custom REST App

** Only statistically significant correlations are reported: |r| > 0.3, p < 0.05*

Results

- **# SQL statements**: r = 0.366 - 0.460, p < 0.006
- **Normal workload** - **CPU utilization** : r = 0.717, p < 0.001
- **Stress & concurrent workloads**: **# HTTP requests, # Network packets, Total packet length**: r > 0.4, $p \leq$ 0.005

Table 4: Pearson correlations between collected metrics and energy consumption for Custom REST App.

	Exp. (i)	Exp. (ii)	Exp. (iii)
Collected metrics	r	r	r
CPU Utilization	0.717***	0.438	0.668
#HTTP Requests	0.030	0.423**	0.461***
#Req. by Method PUT	-	0.423**	0.461***
#Req. by Response 200	0.030	0.423**	0.461***
#SQL statements	0.366***	0.421**	0.460***
#Network Packets	-0.096	0.425**	0.449***
Total Length of Packets	-0.10	0.425**	0.454***

(i) Normal workload; (ii) Stress workload; (iii) Concurrent workload
r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS – Custom REST App

** Only statistically significant correlations are reported: $|r| > 0.3$, $p < 0.05$*

Results

- **# SQL statements**: $r = 0.366 - 0.460$, $p < 0.006$
- **Normal workload - CPU utilization** : $r = 0.717$, $p < 0.001$
- **Stress & concurrent workloads: # HTTP requests, # Network packets, Total packet length**: $r > 0.4$, $p \leq 0.005$

Insights

- **# SQL statements** - the **most stable and reliable** indicator of energy consumption
- **Network and HTTP-related metrics** - *moderately predictive, but mainly under stress or concurrent workloads*
- **CPU utilization** - *strong predictor only under the normal workload*
- **Memory utilization** - *no statistically significant relationship*

Table 4: Pearson correlations between collected metrics and energy consumption for Custom REST App.

	Exp. (i)	Exp. (ii)	Exp. (iii)
Collected metrics	r	r	r
CPU Utilization	0.717***	0.438	0.668
#HTTP Requests	0.030	0.423**	0.461***
#Req. by Method PUT	-	0.423**	0.461***
#Req. by Response 200	0.030	0.423**	0.461***
#SQL statements	0.366***	0.421**	0.460***
#Network Packets	-0.096	0.425**	0.449***
Total Length of Packets	-0.10	0.425**	0.454***

(i) Normal workload; (ii) Stress workload; (iii) Concurrent workload
r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS – Maven

** Only statistically significant correlations are reported: |r| > 0.3, p < 0.05*

Results

- **# Maven logs - activiti**: $r = -0.507$, $p < 0.001$
- **# Maven logs - jackson**: $r = 0.315$, $p < 0.05$
- **Disk reads - in most projects**: $r = 0.306 - 0.514$, $p \leq 0.03$

Table 5: Pearson correlations between collected metrics and energy consumption for Maven projects.

	activiti	clojure	jitwatch	languagetool	litemall	esson	jackson	johnzon	JSON-Java	JSON-Simple
Collected data	*r*	*r*	*r*	*r*	*r*	*r*	*r*	*r*	*r*	*r*
#Maven logs	-0.507***	-0.183	-0.522	-0.116	-0.301	0.113	0.315*	-0.129	<0.001	-0.116
#Logs at level INFO	-0.537***	-0.241	-0.344	-0.116	-0.311	0.113	0.318*	-0.111	0.005	-0.116
#Reads on disk	0.199	0.394 **	0.432*	0.306*	0.430*	0.435**	0.260	0.238	0.348*	0.514**

r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS – Maven

** Only statistically significant correlations are reported: |r| > 0.3, p < 0.05*

Results

- **# Maven logs - activiti**: $r = -0.507$, $p < 0.001$
- **# Maven logs - jackson**: $r = 0.315$, $p < 0.05$
- **Disk reads - in most projects**: $r = 0.306 - 0.514$, $p \leq 0.03$

Insights

- **Disk reads** - **most consistent** energy indicator
- **# Maven logs** - *unreliable predictor* across projects

Table 5: Pearson correlations between collected metrics and energy consumption for Maven projects.

	activiti	clojure	jitwatch	languagetool	litemall	esson	jackson	johnzon	JSON-Java	JSON-Simple
Collected data	r	r	r	r	r	r	r	r	r	r
#Maven logs	-0.507***	-0.183	-0.522	-0.116	-0.301	0.113	0.315*	-0.129	<0.001	-0.116
#Logs at level INFO	-0.537***	-0.241	-0.344	-0.116	-0.311	0.113	0.318*	-0.111	0.005	-0.116
#Reads on disk	0.199	0.394 **	0.432*	0.306*	0.430*	0.435**	0.260	0.238	0.348*	0.514**

r: The Pearson correlation coefficient
$*p < 0.05, **p < 0.01, ***p < 0.001$

RESULTS & INSIGHTS

Cross-Application Insights

- **CPU utilization & SQL statements - most reliable predictors**
- **Disk activity**
  - Disk writes matter in synthetic workloads
  - Disk reads matter in Maven
- **HTTP & general logs** - less reliable
- **Workload and app type** affect predictiveness**:**
  - Controlled and simple workloads - consistent
  - Complex/variable workloads - unpredictable
  - Microservices - one service may use most of the energy

RESULTS & INSIGHTS

Cross-Application Insights

- **CPU utilization & SQL statements - most reliable predictors**
- **Disk activity**
  - Disk writes matter in synthetic workloads
  - Disk reads matter in Maven
- **HTTP & general logs** - less reliable
- **Workload and app type** affect predictiveness:
  - Controlled and simple workloads - consistent
  - Complex/variable workloads - unpredictable
  - Microservices - one service may use most of the energy

Research Question:

Can software energy consumption be estimated using software traces and logs?

RESULTS & INSIGHTS

Cross-Application Insights

- **CPU utilization & SQL statements - most reliable predictors**
- **Disk activity**
  - Disk writes matter in synthetic workloads
  - Disk reads matter in Maven
- **HTTP & general logs** - less reliable
- **Workload and app type** affect predictiveness:
  - Controlled and simple workloads - consistent
  - Complex/variable workloads - unpredictable
  - Microservices - one service may use most of the energy

Research Question:

Can software energy consumption be estimated using software traces and logs?

Answer to Research Question:

Yes, using hardware metrics + structured logs

RESULTS & INSIGHTS

Cross-Application Insights

- **CPU utilization & SQL statements - most reliable predictors**
- **Disk activity**
  - Disk writes matter in synthetic workloads
  - Disk reads matter in Maven
- **HTTP & general logs** - less reliable
- **Workload and app type** affect predictiveness:
  - Controlled and simple workloads - consistent
  - Complex/variable workloads - unpredictable
  - Microservices - one service may use most of the energy

Research Question:

Can software energy consumption be estimated using software traces and logs?

Answer to Research Question:

Yes, using hardware metrics + structured logs

Limitations

- Single Machine
- Limited Workload Diversity
- Scope Restricted to Univariate Analysis

CONCLUSION & FUTURE WORK

CONCLUSION & FUTURE WORK

Key Findings:

- **Hardware metrics** (CPU utilization, disk activity): **consistent predictors** of energy consumption
- **Structured software logs** (SQL statements, Docker container logs): **useful**, especially for single-service bottlenecks
- **Unstructured logs** (HTTP requests, Maven logs): **unstable correlations**, less reliable
- **Workload & application structure:** strongly influence predictive energy consumption

CONCLUSION & FUTURE WORK

Key Findings:

- **Hardware metrics** (CPU utilization, disk activity): **consistent predictors** of energy consumption
- **Structured software logs** (SQL statements, Docker container logs): **useful**, especially for single-service bottlenecks
- **Unstructured logs** (HTTP requests, Maven logs): **unstable correlations**, less reliable
- **Workload & application structure:** strongly influence predictive energy consumption

Future Work:

- Expand experiments to **larger-scale**, **complex workloads**, and **diverse architectures**
- Develop **tools/frameworks** integrating **hardware + software metrics** for **real-time energy monitoring**
- Design **energy-aware logging mechanisms** to improve **prediction and proactive management**

Thank you for listening!

I'm happy to take your questions.

Or reach out to
thi-mai-phuong.nguyen@univ-pau.fr
for any other questions.

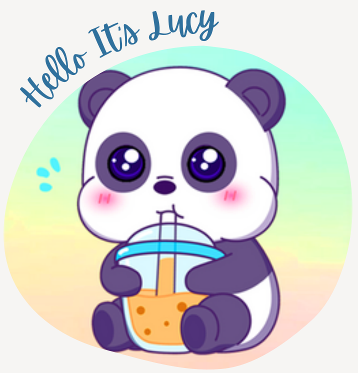

BDCAT 2025

IEEE/ACM International Conference on
Big Data Computing,
Applications and Technologies

UNIVERSITÉ DE PAU ET DES PAYS DE L'ADOUR · LIUPPA · Université Paris Nanterre · LIP6

This research was supported by CCLO (France) and the RESILINK project (PRIMA S2 2021, EU-funded).