Ichnos: A Carbon Footprint Estimator for Scientific Workflows

Kathleen West University of Glasgow Glasgow, United Kingdom kathleen.west@glasgow.ac.uk

Yehia Elkhatib University of Glasgow Glasgow, United Kingdom yehia.elkhatib@glasgow.ac.uk

Abstract

Scientific workflows facilitate the automation of data analysis, and are used to process increasing amounts of data. Therefore, they tend to be resource-intensive and long-running, leading to significant energy consumption and carbon emissions. With ever-increasing emissions from the ICT sector, it is crucial to quantify and understand the carbon footprint of scientific workflows. However, existing tooling requires significant effort from users – such as setting up power monitoring before executing workloads, or translating monitored metrics into the carbon footprints post-execution.

In this paper, we introduce a system to estimate the carbon footprint of Nextflow scientific workflows that enables post-hoc estimation based on existing workflow traces, power models for computational resources utilised, and carbon intensity data aligned with the execution time. We discuss our automated power modelling approach, and compare it with commonly used estimation methodologies. Furthermore, we exemplify several potential use cases and evaluate our energy consumption estimation approach, finding its estimation error to be between 3.9–10.3%, outperforming both baseline methodologies.

1 Introduction

Scientists in many fields, including genomics, materials science, and remote sensing, need to analyse increasing amounts of data [4, 12, 21, 24]. Scientific workflow systems facilitate the automation of such analyses, enabling scientists to compose pipelines out of black-box tasks with data dependencies between them. Because these workflows are often used to process large quantities of data, they tend to be resource-intensive and long-running, leading to significant energy consumption and, therefore, carbon emissions. Furthermore, the growing popularity of big data applications has been identified as a driver of the increasing emissions of the ICT sector [13]. As such, it is crucial to quantify and understand the carbon footprint of scientific workflows.

Scientific workflow systems such as Nextflow [8] allow for the design, execution, and monitoring of workflows on heterogeneous clusters. While these systems usually generate detailed performance traces and logs for executed workflows, they do not produce a record of the energy consumed or carbon emitted. Consequently, users must manually monitor power consumption with hardware/software power meters or, otherwise, use a methodology like Cloud

Magnus Reid

University of Glasgow Glasgow, United Kingdom 2668261r@student.gla.ac.uk

Lauritz Thamsen University of Glasgow Glasgow, United Kingdom lauritz.thamsen@glasgow.ac.uk

Carbon Footprint (CCF)¹ or Green Algorithms (GA) [19], which employ linear power models to translate resource utilisation into energy consumption. In either case, to translate the energy consumed into carbon emitted, users need a measure of carbon intensity (CI), such as a yearly average or a more fine-grained metric. Generally, CI measures the amount of carbon (CO_2e) produced per kilowatthour (kWh) of electricity consumed, and varies across different locations, seasons, and times, depending on the sources generating electricity and the demand on the grid.

In practice, monitoring power consumption requires the user to attach a physical power meter or to enable a software-based tool like Intel's Running Average Power Limit (RAPL) prior to executing a workflow. Without this step, power consumption can only be estimated based on coarse-grained resource utilisation averages. This is possible using the CCF and GA methodologies, but only at reduced accuracy. The GA methodology relies on vendor-specified Thermal Design Power (TDP) of assigned compute resources, a proprietary metric that does not reflect key processor settings, such as processor frequency, and does not indicate idle power consumption. The CCF methodology builds a linear power model between the power consumption measured at 0% and 100%; however, this does not allow for the potential of non-linear increases in power consumption over this range [18]. Furthermore, while both methodologies translate power consumption into carbon emissions, they use a static average value to represent the CI of electricity consumed by the compute workload, ignoring the substantial temporal variability in CI.

To address these limitations, we propose *Ichnos*, a novel and flexible system for estimating the carbon footprint of Nextflow workflows based on detailed workflow traces, resource-specific power models, and CI time-series data. First, Ichnos takes as input the automatically generated workflow trace produced by Nextflow. The use of these traces is an original contribution, ensuring that users do not need to manually monitor power consumption and enabling the analysis of previously executed workflows. Next, Ichnos enables users to automatically generate a power model for utilised compute resources to accurately reflect processor settings, such as processor frequency, instead of solely relying on a linear function between min/max power consumption values per processor – though this is offered as a fallback methodology if users no longer have access to compute resources to generate power models. Finally, Ichnos converts the estimated energy consumption to

LOCO '24, 3rd December, 2024, Glasgow, Scotland 2025.

¹ https://www.cloudcarbonfootprint.org/docs/methodology/

overall carbon emissions using fine-grained time-series CI data for each workflow task and only resorts to coarse-grained yearly averages where high-resolution location-based CI data are not available. Additionally, Ichnos reports estimated energy consumption and carbon emissions per workflow task, providing greater granularity than existing methodologies, and allows users to identify which of their tasks have the largest footprint to address. We provide the implementation of Ichnos as open-source². We evaluate the accuracy of the automated power models generated with Ichnos used to estimate energy consumption and compare them with the monitored energy consumption and baseline methodology estimations from CCF and GA. We demonstrate our estimator system on traces from three real-world Nextflow workflows and show the system's functionality by varying the granularity of provided CI data, using both average and marginal CI, and varying processor governor settings of assigned compute resources.

2 Background

Here, we summarize scientific workflows and carbon intensity.

2.1 Scientific Workflows

Scientific workflows automate data analysis processes that support a scientific objective. They are often defined in terms of their tasks and data dependencies, and are represented using directed acyclic graphs, or as pipelines.



Figure 1: Simple representation of a scientific workflow, composed of seven tasks (A–G).

Figure 1 shows a simple representation of a scientific workflow, with seven tasks (A–G) and the dependencies between them. The tasks that form a workflow are considered black-box processes, receiving input from previous task(s), undergoing some processing and producing output to send on to subsequent tasks. As tasks are individual processes, they could run in parallel on available compute devices. For example, Tasks B and C could run in parallel after receiving input from Task A. Dedicated resource usage data could therefore be generated for each task, such as its runtime, status and the compute resource it was executed on.

Scientific Workflow Management Systems (SWMS) like Kepler, Nextflow, and Pegasus allow scientists to design, execute, and monitor workflows on heterogeneous infrastructure. Such systems enable the generation of performance traces, at a task level.

2.2 Carbon Intensity Signals

The carbon intensity (CI) of electricity measures carbon dioxide released per unit of energy produced. Generating energy from renewable energy sources such as wind or solar reduces CI. CI also typically decreases during periods of time when the demand on the grid is low. In this work, we use grams of carbon dioxide per kilowatt-hour (qCO2e/kWh) to measure CI.

CI can be quantified using two signals: average and marginal. Average reflects the overall grid emissions at the time when electricity is requested, factoring in each energy source's relative share and emission rate. Marginal measures the emissions of the specific energy source(s) used to meet additional load at the time when electricity is requested. While marginal CI is preferred for measuring the impact of demand shifting, the average CI is more readily available and used for emissions accounting. Moreover, average and marginal CI are available at different levels of granularity, such as yearly averages, hourly values, and 5/15/30-minute values. As a general-purpose system estimating the carbon footprint, we leave the CI signal choice to the user to provide the greatest flexibility.

3 System Design

In this section, we discuss the requirements for the design of Ichnos and also provide an overview of the estimator system's design.

3.1 Requirements

We identify the following requirements from which we derive the design of Ichnos.

Enable post-hoc estimation. We enable post-hoc use of Ichnos after workflows have been executed. This enables users to analyse the carbon footprint of previous **and** new experiments, from executions that may have occurred on local devices, clusters, and cloud infrastructure. Given this, users may no longer have access to the infrastructure used to execute workflows, whether the devices were replaced, or only temporarily booked in a cloud environment.

Use resource utilisation data. Resource utilisation monitoring is often more readily available than power monitoring. Users typically lack access to power monitoring tools in public cloud environments. Also, users may not have configured power monitoring when workflows were executed, but will often still have access to monitoring data. Hence, Ichnos uses existing workflow traces that contain task-level resource usage data, without energy measurements.

Estimate CPU and Memory Energy Consumption. We focus on the energy consumption of CPU and memory as this typically has the largest dynamic power consumption range to attribute to specific load on compute resources. In addition, workflows are often executed on shared resources, where storage access may not be limited to the nodes a workflow's tasks execute on. However, workflow traces *are* commonly limited to resource usage by tasks on specific nodes.

Estimate Operational Carbon Emissions. We currently estimate the operational carbon emissions. While electricity grid emissions data, such as average and marginal CI, are increasingly available, Life Cycle Assessments (LCAs) detailing embodied carbon emissions are less widely available for specific server hardware. Furthermore, for virtual resources in public cloud infrastructure,

²https://github.com/GlasgowC3lab/ichnos

we do not have detailed information on the utilised server hardware.

Estimate Workflow Carbon Footprint. Nextflow workflow traces are produced for individual workflow runs, and offer no insight into other processes running on the same machines. We, therefore, consider the estimation of the overall system carbon footprint in multi-tenant compute scenarios, such as public cloud infrastructure, to be outside the scope of our system, and rely on the generated workflow traces for accounting for the correct shares of an overall system.

3.2 System Overview

Ichnos is a system that produces an estimate of the operational carbon footprint from the execution trace of a Nextflow scientific workflow using power and energy data aligned with the execution. Figure 2 provides an overview of the design.



Figure 2: High-level design of the Ichnos Carbon Footprint estimator system with per-task power and emissions estimation, based on provided input data, and detailed reporting.

In Phase 1, the user provides input in the form of three items:

- The workflow trace that includes a task-level summary of resource usage including the runtime, CPU utilisation, and allocated memory;
- (2) The power model selected to estimate the power consumption to accurately estimate energy consumption, the user can select an automatically generated power function or regression-based model to reflect processor settings; and
- (3) The CI data which should be fine-grained time-series data, if available, or instead a coarse-grained average.

In Phase 2, resource usage data are extracted from the workflow trace for each task, and the energy consumption is estimated using the selected power model. Subsequently, the energy consumption per task is translated into carbon emissions using the provided CI data. This estimates operational carbon emissions by aligning the tasks of potentially long-running workflow applications with CI data matching the specific execution times. These estimations are summed to calculate the power consumption and carbon emissions for the overall workflow execution. In Phase 3, the energy consumption and carbon emissions estimated for each task are summarised in a trace file, alongside a summary of the overall carbon footprint. We also produce a trace file identifying the 10 most energy-intensive and the 10 longestrunning workflow tasks, allowing users to review their relative emissions, and to consider the potential of aligning tasks with fluctuating CI – applying existing carbon-aware computing methods to reduce the overall footprint [15, 27].

3.3 Automated Power Modelling

Ichnos supports generating power models for utilised compute resources in an automated manner - selecting the most accurate available model. The estimator system does this in a two-step procedure, first taking several measurements of energy consumed by CPU and memory by stress-testing the CPU and memory using the Turbostress³ tool. Ichnos has a default setting of 11 measurements, adjusting CPU load from 0% (idle) to 100% in 10% increments. These measurements should ideally be taken at the time of execution, on the compute resources where workflows were executed, documenting processor settings such as the governor selected (which decides how the CPU frequency is adjusted based on CPU demand), and the date to version readings. The second step involves generating a power model from the readings and fitting a polynomial or linear model from these readings. The power modelling phase should be run at regular intervals through compute resource lifetimes as well as when hardware is changed, to account for altered device performance. We evaluate the accuracy of Ichnos's power model generation in Section 4.2.

In the scenario where users estimate energy consumed by a historical workflow, executed on resources that they no longer have access to, or on public cloud resources that have been released, or anywhere that a user cannot execute the energy measurements script - Ichnos has the fallback option of using a linear power model, or if only the CPU model is known, a per-core value based on vendor-specified TDP. Both these fallbacks are used in existing estimation methodologies and can offer ballpark estimates of the energy consumption to then translate into carbon emissions. The memory energy consumption coefficient in Ichnos can be configured according to the Turbostress readings. Otherwise, the tool defaults to a constant conversion factor, as used in existing estimation methodologies. The memory stress testing is less comprehensive than that of the CPU. However, we found that memory energy consumption generally has a relatively small impact on the overall workflow energy consumption.

3.4 Carbon Footprint Reporting

Ichnos offers human-readable files to quickly understand the footprint of a workflow for general reporting. It also produces computerreadable files as an augmented trace file where the tasks' carbon emissions and energy consumption are reported. This is provided to enable scientists to better understand their workflow's footprint, and the heavy-hitting tasks that may be disproportionately contributing to the overall footprint.

The following three files are generated:

³https://github.com/teads/turbostress

LOCO '24, 3rd December, 2024, Glasgow, Scotland

name	id	co2e	energy	avg_ci	realtime	cores	usage (%)
DADA2_ERR	58	2.742	0.044	55.0	2576181	6	306.0
PICRUST	66	2.337	0.034	60.5	1977266	6	340.0
DADA2_DENOISING	59	0.678	0.011	56.0	651643	6	276.0
BARRNAP	65	0.347	0.006	56.0	348293	2	198.0
DADA2_ADDSPECIES	67	0.285	0.005	56.0	288655	1	100.0
DADA2_TAXONOMY	64	0.278	0.005	56.0	204224	16	1310.0
DADA2_RMCHIMERA	60	0.074	0.001	56.0	66396	6	506.0
DADA2_QUALITY1	40	0.054	0.001	49.0	63627	2	102.0
DADA2_QUALITY1	39	0.054	0.001	49.0	63645	2	101.0
MULTIQC	70	0.052	0.001	65.0	48646	1	57.0

Table 1: Extract from detailed summary file listing tasks with the highest carbon footprint and energy consumption.

- a human-readable summary of the workflow, detailing parameters provided and overall energy consumption and carbon emissions;
- (2) a computer-readable file with the carbon emissions and energy consumption estimations of individual tasks;
- (3) a computer-readable summary with the top 10 tasks that had the highest carbon footprint and energy consumption.

Table 1 shows an extract from the summary listing the top 10 tasks with the highest carbon footprints and energy consumption. In the example, we can identify tasks that contribute to the overall footprint, with, for example, PICRUST consuming 0.034kWh, almost 6x as much as BARRNAP, and, since both are executed on similar carbon-intensive energy, a similarly larger carbon footprint.

4 Evaluation

In this section, we first present the experimental setup that we used to evaluate Ichnos. We then analyse the accuracy of the estimator system's generated power models for estimating power consumed during workflow execution, and demonstrate a series of practical use cases for Ichnos.

4.1 Experimental Setup

We use a set of baselines, workflows, and infrastructure for multiple experiments as described here.

Baselines. To evaluate the accuracy of our approach when estimating energy consumption – and therefore carbon emissions –, we compare it against two baseline methodologies that also enable post-hoc footprint estimation: Naive–Linear and Green Algorithms (GA). Naive–Linear is aligned with the Cloud Carbon Footprint (CCF) estimation methodology¹, assuming that energy consumption scales linearly between the idle and maximum power consumption reported by manufacturers or online databases like SPEC⁴. GA uses the manufacturer-specified processor TDP to estimate percore energy consumption. In addition to comparing against both baselines, we compare all estimations against power consumption readings taken with the Perf tool, which uses Intel's RAPL, which we consider to be the ground truth in our experiments.

Workflows. To evaluate our system, we selected two real-world workflows from the nf-core repository⁵ – a community-curated collection of workflows implemented using Nextflow [10]. We chose

the AmpliSeq and NanoSeq workflows, both are bioinformatics workflows within the top 10 most popular workflows from nf-core. We manually executed these workflows to produce experimental data with energy consumption monitoring. In addition, we exemplify the system's capacity for post-hoc carbon footprint estimation by using historical traces where the Rangeland workflow was executed on an infrastructure we do not have access to.

Infrastructure. For the following power modelling accuracy analysis and two of the use cases, we used five nodes from a local heterogeneous cluster – these nodes each have the prefix gpgnode. Within the cluster, gpgnodes 13–16 are each equipped with two Intel Xeon E5-2640 v2 processors, which have a 1.2-2.5GHz frequency range, and 64 GB of RAM. In contrast, gpgnode 22 is equipped with an Intel Xeon Gold 6426Y processor, which has a frequency range of 0.8-4.1GHz, and 128 GB of RAM. Meanwhile, we describe the infrastructures relevant for the other use case demonstrations in the respective subsections.

4.2 Power Modelling Accuracy

We discuss the accuracy of Ichnos's generated power models, in relation to the actual energy consumption, before comparing our estimations with the selected baselines of Naive–Linear and GA.

Power Model Accuracy. For each available compute device, we took power consumption readings, as detailed in Section 3.3. We used these readings to generate cubic and linear regression models of consumed energy.

In Figure 3, we show a plot comparing power models generated for three compute nodes (gpgnode 15, 16, and 22), using the governors: performance, powersave, and ondemand – where available. The original Turbostress readings are marked on the plots. To these readings, Ichnos fits a cubic model (Ichnos–Cubic) and a linear one (Ichnos–Linear). We additionally plot the Naive–Linear model, which assumes linear scaling of power consumption from the readings at 0% (idle) and 100% (peak) utilisation.

Despite these nodes being equipped with the same resources, their peak power consumption can vary substantially. For example, gpgnode-15 using the performance governor reached ~140*W* at 100% utilisation while gpgnode-16 reported a value ~110*W* at the same utilisation, using the same governor settings. This difference (\approx 27%) underscores the importance of resource-specific power modelling based on actual measurements. The selected governor also impacts the power model, with the powersave governor consistently using less power at 100% utilisation.

To further demonstrate how hardware generation affects power consumption patterns, readings for gpgnode-22 are plotted. This newer hardware not only consumes significantly more power as load is increased from 0-100%, but also exhibits markedly different behaviour. We observe the non-linear nature of the power consumed in this node in particular, which highlights why linear models can be inadequate. This non-linearity is especially pronounced in the 60-80% utilisation range for this node, where the power consumption does not increase significantly with utilisation.

We also report the Root Mean Square Error (RMSE) between the model predicted values and the power consumption readings taken using Turbostress (which uses RAPL to take measurements), for

⁴https://www.spec.org/power_ssj2008/

⁵https://github.com/nf-core



Figure 3: Comparison of power consumption readings over relative load for three compute nodes utilising different Intel governors (as available on processors).

Nada	Corroran	Naine Lincon	Ichnos		
Node	Governor	Naive Linear	Linear	Cubic	
gpgnode-13	ondemand	12.96	6.35	0.59	
"	performance	12.14	5.83	0.84	
"	powersave	7.95	4.02	0.22	
gpgnode-14	ondemand	12.21	5.97	0.61	
"	performance	11.55	5.49	0.86	
"	powersave	7.45	3.77	0.29	
gpgnode-15	ondemand	12.96	6.36	0.62	
"	performance	12.01	5.72	0.85	
"	powersave	8.82	4.43	0.48	
gpgnode-16	ondemand	10.33	5.13	0.34	
"	performance	10.10	4.88	0.53	
"	powersave	6.94	3.52	0.47	
gpgnode-22	performance	56.45	28.78	9.61	
"	powersave	63.10	32.34	10.60	

 Table 2: The RMSE of selected models energy consumption
 estimated executing the Ampliseq workflow.

Ichnos–Cubic, Ichnos–Linear and Naive–Linear models. These are listed in Table 2. Across all selected nodes, Ichnos–Cubic demonstrated superior accuracy, markedly outperforming Ichnos–Linear which, in turn, proved more accurate than the Naive–Linear model.

Use of Generated Power Models to Estimate Energy Consumption. We executed the Ampliseq workflow on gpgnodes 13–16 and 22, on selected governors and monitored the energy consumed. We used Ichnos to generate estimations from the workflow traces using the Ichnos–Cubic, Ichnos–Linear, Naive–Linear, and GA models.

For each node using each specified processor governor setting, we executed the workflow 3 times and report the mean energy consumption. We compare the estimated energy consumption, with the actual consumption recorded using Perf, and report the overall percentage error in Table 3. We found that Ichnos–Linear performed the best for all our estimations, in line with our expectations. Discussion of the Experimental Accuracy of Energy Estimation. In our evaluation of the accuracy of generated power models, Ichnos– Cubic was found to be the most accurate – reporting the minimum RMSE between model predictions and Turbostress readings. However, when we used the power models to estimate energy consumption for real Nextflow executions and reported the percentage error between the estimated and actual energy consumption, we found that Ichnos–Linear consistently outperformed Ichnos–Cubic.

Given these results, we identify the following two drawbacks of using a non-linear model like Ichnos–Cubic to estimate energy consumption.

To enable trace-based resource estimations, we estimate the energy consumption for all individual tasks, summing each task's share of utilised resources. This works when we use a linear model for several tasks running in parallel on the same shared compute resource. However, if we use a cubic model and consider a task that has an average CPU utilisation of 100% on one core, its power consumption will differ depending on whether it runs on its own or shares resources with other CPU-intensive tasks. This effect can be seen in Figure 4, which shows the power consumption of the commonly used bioinformatics task FastQC⁶ at different system background loads. This task fully utilises one core on each system. We can see that the power consumption is significantly affected by the background system load, and this is especially notable for gpgnode 22, which shows a large difference between 0–10% utilisation. This result is expected, and reflects the behaviour seen in Figure 3.

Furthermore, the traces generated from Nextflow executions only provide coarse-grained CPU utilisation averages for each individual workflow task – which could have a runtime spanning seconds to hours in length and markedly different utilisation over time.

Therefore, we recommend use of the generated Ichnos–Linear model when estimating energy consumption to reduce the power modelling estimation error – compared to Naive–Linear and GA – and avoid the estimation being affected by background task utilisation and how resource utilisation data are aggregated.

⁶https://github.com/s-andrews/FastQC

Nada	Governor	Perf	Ichnos-Cubic	Error	Ichnos-Linear	Error	Naive-Linear	Error	GA	Error
node		(kWh)	(kWh)	(%)	(kWh)	(%)	(kWh)	(%)	(kWh)	(%)
gpgnode-13	ondemand	0.161	0.135	16.1	0.144	10.3	0.121	24.7	0.28	82.9
gpgnode-14	performance	0.161	0.138	14.2	0.146	9.1	0.124	22.8	0.026	83.7
"	powersave	0.159	0.143	9.8	0.150	5.6	0.136	14.4	0.029	81.4
gpgnode-15	performance	0.168	0.147	12.4	0.155	7.4	0.134	19.9	0.027	83.6
"	powersave	0.178	0.157	11.7	0.165	7.3	0.148	16.7	0.031	82.4
gpgnode-16	ondemand	0.139	0.124	10.8	0.131	5.4	0.113	18.8	0.026	81.4
gpgnode-22	performance	0.165	0.131	20.7	0.159	3.9	0.101	38.7	0.003	98.0
"	powersave	0.163	0.031	81.0	0.150	8.0	0.085	47.9	0.003	98.0

Table 3: The estimated energy consumption using generated models with percentage error.



Figure 4: Comparison over two nodes of power consumed by a task as background system load varies.

4.3 Ichnos Use Cases

Use of Historical Traces. We used Ichnos to estimate the carbon footprint of historical executions reported for the Rangeland workflow. Specifically, we estimated the carbon footprint of the three workflow executions that occurred on a single node in Germany equipped with an Intel Xeon Silver 4314 processor and 256 GB of RAM. We configured the system to use the yearly average CI in Germany in 2023 as $394gCO_2e/kWh$. As we did not have access to perform power measurements on the node at the time of execution, we used a linear power model ranging between 80-135W to estimate energy consumption. The average energy consumption was **30.51***kWh*, with CPU energy consumption accounting for 99% of the overall energy consumption and memory responsible for the remaining 1%. The estimated carbon footprint was **12***kgCO*₂*e*.

Use of Cloud Traces. We executed the NanoSeq workflow using the Google Cloud Batch execution environment with Nextflow. The workflow ran on n2 and c2 cloud nodes in the europe-west-2 region. We had access to the workflow trace generated, but did not have access to the utilised nodes to retrieve energy measurements to fit a power model. Such a scenario may resemble scientists' realworld use cases for Ichnos, where they have limited information on shared nodes and also on the specific location of these to access accurate CI data. Despite these challenges, Ichnos allows for posthoc estimation.



Figure 5: Variable Carbon Intensity in South Scotland on 26/09/2024, with workflow execution highlighted in green.

We configured the system to use a linear power model, and made use of CCF's coefficients repository⁷ to obtain measurements for the Cascade Lake machine family – used by the n2 and c2 machine types. These coefficients stated a minimum of 0.69W per-core, and a maximum of 3.75W per-core. We adjusted these values in line with the number of cores requested for each task in our estimations. We used the quoted average CI value for the europe-west-2 region that Google puts⁸ at 136gCO2e/kWh. Using Ichnos, we estimated the carbon emissions to be 15.5gCO2e, translated from an estimated energy consumption of 0.114kWh.

Varying the Granularity of CI Data. We took the workflow trace from the AmpliSeq execution that ran on the edge server, equipped with an Intel i7-10700T processor with 32 GB of RAM, for 2h 40m in the evening of September 26th 2024 in Glasgow. The CI fluctuation for South Scotland region of the National Grid⁹ is depicted in Figure 5. If this information was not available, we would use the average CI for the National Grid in 2023, which was 215*gCO2e*¹⁰.

Ichnos estimated that the footprint was **0.33***g*CO2*e* using the region-specific time-series CI data, whereas the footprint estimated using the coarse-grained average would be **18.65***g*CO2*e* - an estimate almost 60x larger. This highlights the potential of using a flexible system, such as Ichnos, where the user can provide specific high-resolution time-series CI data to estimate the carbon footprint more accurately.

Varying Processor Governor Settings. We executed the AmpliSeq workflow using the full-size real-world dataset provided from nf-core on gpgnodes 13–16 and 22. Each compute node used

⁷https://github.com/cloud-carbon-footprint/ccf-coefficients

⁸https://cloud.google.com/sustainability/region-carbon

⁹https://carbonintensity.org.uk/

¹⁰https://app.electricitymaps.com/

2–3 available governors, considering the real-world scenario where users may not be in control of the governor in use, and cannot set a node to a fixed frequency. We monitored the energy consumption using Perf again for comparison. We estimated the carbon emissions using Ichnos configured with the fitted linear model (Ichnos–Linear) as discussed in Section 4.2. We compare the estimations with the Naive–Linear and GA baselines.

Table 4 lists the experimental results. We cannot directly evaluate the accuracy of our carbon footprint estimations, as Perf only monitors energy consumption, rather than carbon emissions. However, the accuracy of the selected power models can be considered. We found that energy consumption estimations from Ichnos–Linear had an error of 3.9–10.3%, compared to Naive–Linear's error of 14.4–47.9%, and GA's error of 81.4–98.0%. Ichnos translated the estimated energy consumption for each of these power models using the same CI data, at the same level of granularity, carrying any estimation error linearly to the carbon footprint estimates. Therefore, Ichnos enables post-hoc estimation with a reduced estimation error, outperforming both baselines.

Use of Average and Marginal CI Data. Ichnos offers its users the opportunity to provide CI data in an input file of dynamic intervals of granularity. This also allows users to provide average or marginal CI data, which are often available at different levels of granularity, from different providers.

To demonstrate Ichnos's flexibility, we show the footprint estimations for executions of the Ampliseq workflow on gpgnode-13–16 and 22, where the available governors were used. These executions occurred in February 2025, and we retrieved CI data aligning with the actual execution times. Average CI data were available for 30-minute intervals from the National Grid for the South Scotland region where the servers were located. Marginal CI data was available for 5-minute intervals from WattTime¹¹ for the United Kingdom, yet more regional data was not available.

The estimated carbon emissions produced by Ichnos – using the Ichnos–Linear power model – are shown in Table 5, demonstrating that the system enables footprint estimation with the most granular CI data available, offering users the choice of signal.

5 Limitations

Generality of Ichnos. Ichnos was created to estimate the carbon footprint of Nextflow workflow executions from automatically generated trace files. The system is configured with the expected format of this trace file, constructing records for each individual task. Therefore, it is not directly applicable to other workflow systems. However, the method used to estimate energy consumption, and to translate such consumption into carbon emissions using CI data could be adapted to work with performance trace file formats used for other workflow systems, like Pegasus or Makeflow. We welcome contributions to the open-source estimator system, which remains under active development.

Power Model Generation and Use. Ichnos allows users to run a series of energy consumption readings by varying the computational load on a CPU, generating a resource-specific power model from these readings for use when estimating the carbon footprint. This power model generation process should be repeated when the hardware is changed or at regular intervals to track device degradation. However, this is reliant on the user having the correct permissions available and access to compute infrastructure on which workflows are executed. If power models are not up-todate or do not align with the original workflow execution times, the accuracy of the energy consumption estimations will be limited.

Limitations of Input Data. Ichnos allows users to provide CI data at varied levels of granularity – enabling the use of both average and marginal CI. However, CI data usually specify a value over a given period of time, e.g. WattTime offers marginal CI at intervals of five-minutes, while the National Grid offers average CI at intervals of thirty-minutes. As these intervals become less granular, the overall footprint estimation becomes less accurate. Furthermore, we are reliant on these CI data sources for supplying accurate data.

By enabling post-hoc estimation, where the user will likely not have had access to power or emissions monitoring tools, we can only guarantee that our tool uses the estimation methodology described and the user-provided data – the workflow trace and the CI as well as the power model generated from measurements on compute resources utilised at the time, or as close as possible. This limitation is the same for other existing footprint estimation methodologies like CCF and GA.

6 Related Work

This section examines energy consumption measurement methods in a broad context, followed by carbon footprint estimation methodologies and prior research wherein the carbon footprint of scientific workflows has been explored.

To estimate the carbon footprint of computation, power consumption must first be monitored with power meters, or measured with energy profiling tools. These traditionally rely on software interfaces like the Running Average Power Limit (RAPL) – available on Intel Processors – or the NVIDIA Management Library (NVML) – available on NVIDIA GPUs. Tools built using these interfaces, such as Nvidia-smi, Perf and Scaphandre, can provide accurate measurements of energy consumption [17]. However, these necessitate configuration prior to workload execution. Given the requirement that our estimations are made post-hoc, these are not suitable in our problem space. We instead focus on methods that are capable of modelling power consumption based on compute resource usage.

Many methods have been proposed to model the power consumption at server and data-centre levels [1]. Some works consider the power consumption of a server to be the sum of idle consumption – thought to be a fixed value – and active consumption caused by computational workloads [7, 22, 28]. Other studies use regression models to predict power consumption based on server properties in combination with the idle power consumption [9, 11, 20, 29]. Some works consider the CPU utilisation to be the dominant contributor when modelling server consumption (for example, the linear model given by Fan et al. [11]) which produced reasonably accurate estimations and has since been implemented in various estimation

¹¹ https://watttime.org/

Node	Governor	Perf	Ichnos-Linear	Emissions	Naive-Linear	Emissions	GA	Emissions
Node		(kWh)	(kWh)	(gCO2e)	(kWh)	(gCO2e)	(kWh)	(gCO2e)
gpgnode-13	ondemand	0.161	0.144	2.93	0.121	2.46	0.28	0.56
gpgnode-14	performance	0.161	0.146	2.96	0.124	2.51	0.026	0.53
"	powersave	0.159	0.150	1.78	0.136	1.61	0.029	0.36
gpgnode-15	performance	0.168	0.155	1.79	0.134	1.54	0.027	0.31
"	powersave	0.178	0.165	4.05	0.148	3.64	0.031	0.76
gpgnode-16	ondemand	0.139	0.131	7.89	0.113	6.77	0.026	1.55
gpgnode-22	performance	0.165	0.159	6.52	0.101	4.16	0.003	0.13
"	powersave	0.163	0.150	8.35	0.085	4.73	0.003	0.18

Table 4: The estimated energy consumption and carbon emissions using Ichnos-Linear and baseline models.

Table 5: The estimated energy consumption and carbon emissions using average and marginal CI data.

Nada	Governor	Energy	Emissions (gCO2e)			
Noue		(kWh)	Average	Marginal		
gpgnode-13	ondemand	0.144	2.93	141.95		
gpgnode-14	performance	0.146	2.96	142.94		
"	powersave	0.150	1.78	154.53		
gpgnode-15	performance	0.155	1.79	159.30		
"	powersave	0.165	4.05	157.93		
gpgnode-16	ondemand	0.131	7.89	138.75		
gpgnode-22	performance	0.159	6.52	149.92		
"	powersave	0.150	8.35	156.08		

methodologies, including CCF. From the many models proposed, we focused on implementing a fitted linear-regression model, a fitted cubic-regression model and compared these with models that are commonly used by existing estimation methodologies.

Several tools have been created to estimate the carbon footprint from computational workloads [2, 5, 6, 16, 26]. Many of these tools model power consumption by using the server utilisation and the TDP reported by the manufacturer [5, 6, 19]. However, this value does not consider processor settings such as the frequency, reducing the estimation's accuracy. Other tools require the user to have privileged (root) access, and to configure tools prior to workload execution [2, 6, 16, 26]. With our requirement of enabling post-hoc estimation, these tools are out of scope. We focus on existing carbon footprint estimation methodologies, which also enable post-hoc estimation and involve an intermediate step where energy consumption is estimated: Cloud Carbon Footprint¹ (CCF) and Green Algorithms [19] (GA). We compare Ichnos against both methodologies in our evaluation.

Some research studies have specifically applied existing carbon footprint estimation methodologies to analyse the footprint of bioinformatics [14] and neuroimaging [25] research processes. In other works where the focus is instead on reducing the energy footprint of scientific workflows, linear power models have been employed to estimate power consumption [23]. The presence of these works clearly indicates interest in being able to estimate the carbon footprint of computation, and validate the use of estimation methodologies in a post-hoc manner.

7 Conclusion

In this paper, we presented Ichnos – a new system to estimate the carbon footprint of Nextflow workflow executions based on a given trace file, generated power models, and configured carbon intensity data. The system enables the post-hoc estimation of energy and carbon footprints. To improve the accuracy of the estimated power consumption, Ichnos allows for a series of power measurements to be taken to create a power model for estimating the energy consumption. These measurements are repeatable, enabling users to update the models when processor settings change or device performance degrades.

We evaluated the accuracy of the generated power models in comparison to existing carbon footprint estimation methodologies. We showed that Ichnos reduced the estimation error to between 3.9– 10.3%, compared to Naive–Linear's 14.4–47.9% and GA's error of 81.4–98.0%. We further exemplified the estimator system's use cases, where users could estimate the footprint from historical workflow executions, or provide as granular CI data as they have access to (ranging, for example, from a yearly to five-minute resolution), enabling users to consider both average and marginal CI data.

In the future, we plan to test our method for workflow executions on distributed, heterogeneous clusters with node-specific, versioned power models, as we currently use a single power model for workflows executed across multiple nodes. Furthermore, we will add support for the estimation of embodied carbon emissions and are keen to explore integrating Ichnos with methods that predict workflow performance (e.g. [3]).

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council under grant number UKRI154. We also gratefully acknowledge the sources of electricity grid data: NESO Open Data and Electricity Maps historical data for average carbon intensity as well as marginal operating emission rates calculated by WattTime.

Rights Retention

For the purpose of open access, we have applied a Creative Commons Attribution (CC BY) licence to this manuscript.

Data/Code Availability

An open-source implementation of Ichnos is available at https: //github.com/GlasgowC3lab/ichnos. Ichnos: A Carbon Footprint Estimator for Scientific Workflows

LOCO '24, 3rd December, 2024, Glasgow, Scotland

References

- Kazi Main Uddin Ahmed, Math H. J. Bollen, and Manuel Alvarez. 2021. A Review of Data Centers Energy Consumption and Reliability Modeling. *IEEE Access* 9 (2021).
- [2] Lasse F. Wolff Anthony, Benjamin Kanding, and Raghavendra Selvan. 2020. Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models. ICML Workshop on Challenges in Deploying and monitoring Machine Learning Systems.
- [3] Jonathan Bader, Fabian Lehmann, Lauritz Thamsen, Ulf Leser, and Odej Kao. 2024. Lotaru: Locally predicting workflow task runtimes for resource management on heterogeneous infrastructures. *Future Generation Computer Systems* 150 (2024).
- [4] Bruce Berriman, Ewa Deelman, John Good, Joseph Jacob, Daniel Katz, Carl Kesselman, Anastasia Laity, Thomas Prince, Gurmeet Singh, and Mei-Hui Su. 2004. Montage: A Grid Enabled Engine for Delivering Custom Science-Grade Mosaics on Demand. In Optimizing Scientific Return for Astronomy through Information Technologies, Vol. 5493.
- [5] S. A. Budennyy, V. D. Lazarev, N. N. Zakharenko, A. N. Korovin, O. A. Plosskaya, D. V. Dimitrov, V. S. Akhripkin, I. V. Pavlov, I. V. Oseledets, I. S. Barsola, I. V. Egorov, A. A. Kosterina, and L. E. Zhukov. 2023. eco2AI: Carbon Emissions Tracking of Machine Learning Models as the First Step Towards Sustainable AI. Doklady Mathematics (2023).
- [6] Benoit Courty, Victor Schmidt, Sasha Luccioni, Goyal-Kamal, MarionCoutarel, Boris Feld, Jérémy Lecourt, LiamConnell, Amine Saboni, Inimaz, supatomic, Mathilde Léval, Luis Blanche, Alexis Cruveiller, ouminasara, Franklin Zhao, Aditya Joshi, Alexis Bogroff, Hugues de Lavoreille, Niko Laskaris, Edoardo Abati, Douglas Blank, Ziyao Wang, Armin Catovic, Marc Alencon, Michał Stęchły, Christian Bauer, Lucas Otávio N. de Araújo, JPW, and MinervaBooks. 2024. mlco2/codecarbor: v2.4.1.
- [7] Gaurav Dhiman, Kresimir Mihic, and Tajana Rosing. 2010. A system for online power prediction in virtualized environments using gaussian mixture models. In Proceedings of the 47th Design Automation Conference.
- [8] Paolo Di Tommaso, Maria Chatzou, Evan W. Floden, Pablo Prieto Barja, Emilio Palumbo, and Cedric Notredame. 2017. Nextflow Enables Reproducible Computational Workflows. *Nature Biotechnology* 35, 4 (2017).
- [9] Dimitris Economou, Suzanne Rivoire, Christos Kozyrakis, and Parthasarathy Ranganathan. 2006. Full-system power analysis and modeling for server environments. (2006).
- [10] Philip A Ewels, Alexander Peltzer, Sven Fillinger, Harshil Patel, Johannes Alneberg, Andreas Wilm, Maxime Ulysse Garcia, Paolo Di Tommaso, and Sven Nahnsen. 2020. The nf-core framework for community-curated bioinformatics pipelines. *Nature biotechnology* 38, 3 (2020).
- [11] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre Barroso. 2007. Power provisioning for a warehouse-sized computer (ISCA '07).
- [12] James A. Fellows Yates, Thiseas C. Lamnidis, Maxime Borry, Aida Andrades Valtueña, Zandra Fagernäs, Stephen Clayton, Maxime U. Garcia, Judith Neukamm, and Alexander Peltzer. 2021. Reproducible, Portable, and Efficient Ancient Genome Reconstruction with Nf-Core/Eager. *PeerJ* 9 (2021).
- [13] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon S. Blair, and Adrian Friday. 2021. The Real Climate and Transformative Impact of ICT: A Critique of Estimates, Trends, and Regulations. *Patterns* 2, 9 (2021).
- [14] Jason Grealey, Loïc Lannelongue, Woei-Yuh Saw, Jonathan Marten, Guillaume Méric, Sergio Ruiz-Carmona, and Michael Inouye. 2022. The Carbon Footprint of Bioinformatics. *Molecular Biology and Evolution* 39, 3 (2022).

- [15] Walid A Hanafy, Qianlin Liang, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. Carbonscaler: Leveraging cloud workload elasticity for optimizing carbonefficiency. Proceedings of the ACM on Measurement and Analysis of Computing Systems 7, 3 (2023).
- [16] Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning.
- [17] Mathilde Jay, Vladimir Ostapenco, Laurent Lefèvre, Denis Trystram, Anne-Cécile Orgerie, and Benjamin Fichel. 2023. An experimental comparison of softwarebased power meters: focus on CPU and GPU. In 2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid).
- [18] Chaoqiang Jin, Xuelian Bai, Chao Yang, Wangxin Mao, and Xin Xu. 2020. A review of power consumption models of servers in data centers. *Applied Energy* 265 (2020).
- [19] Loïc Lannelongue, Jason Grealey, and Michael Inouye. 2021. Green Algorithms: Quantifying the Carbon Footprint of Computation. Advanced Science 8, 12 (2021).
- [20] Yanfei Li, Ying Wang, Bo Yin, and Lu Guan. 2012. An online power metering model for cloud environment. In 2012 IEEE 11th International Symposium on Network Computing and Applications.
- Paul Muir, Shantao Li, Shaoke Lou, Daifeng Wang, Daniel J. Spakowicz, Leonidas Salichos, Jing Zhang, George M. Weinstock, Farren Isaacs, Joel Rozowsky, and Mark Gerstein. 2016. The Real Cost of Sequencing: Scaling Computation to Keep Pace with Data Generation. *Genome Biology* 17, 1 (2016).
 Swapnoneel Roy, Atri Rudra, and Akshat Verma. 2013. An energy complex-
- [22] Swapnoneel Roy, Atri Rudra, and Akshaf Verma. 2013. An energy complexity model for algorithms. In Proceedings of the 4th conference on Innovations in Theoretical Computer Science.
- [23] Youssef Saadi, Soufiane Jounaidi, Said El Kafhali, and Hicham Zougagh. 2023. Reducing Energy Footprint in Cloud Computing: A Study on the Impact of Clustering Techniques and Scheduling Algorithms for Scientific Workflows. *Computing* 105, 10 (2023).
- [24] Jörg Schaarschmidt, Jie Yuan, Timo Strunk, Ivan Kondov, Sebastiaan Huber, Giovanni Pizzi, Leonid Kahle, Felix Bölle, Ivano Castelli, Tejs Vegge, Felix Hanke, Tilmann Hickel, Jörg Neugebauer, Celso Rêgo, and Wolfgang Wenzel. 2021. Workflow Engineering in Materials Design within the BATTERY 2030 + Project. Advanced Energy Materials 12 (2021).
- [25] Nicholas E Souter, Nikhil Bhagwat, Chris Racey, Reese Wilkinson, Niall W Duncan, Gabrielle Samuel, Loïc Lannelongue, Raghavendra Selvan, and Charlotte L Rae. 2024. Measuring and reducing the carbon footprint of fMRI preprocessing in fMRIPrep. Technical Report.
- [26] Martin Jaggi Tristan Trebaol, Mary-Anne Hartley and Hossein Shokri Ghadikolaei. 2020. A tool to quantify and report the carbon footprint of machine learning computations and communication in academia and healthcare. *Infoscience EPFL:* record 278189 (2020).
- [27] Philipp Wiesner, Ilja Behnke, Dominik Scheinert, Kordian Gontarska, and Lauritz Thamsen. 2021. Let's wait awhile: How temporal workload shifting can reduce carbon emissions in the cloud. In Proceedings of the 22nd International Middleware Conference.
- [28] Peng Xiao, Zhigang Hu, Dongbo Liu, Guofeng Yan, and Xilong Qu. 2013. Virtual machine power measuring technique with bounded error in cloud environments. *Journal of Network and Computer Applications* 36, 2 (2013).
- [29] Xiao Zhang, Jian-Jun Lu, Xiao Qin, and Xiao-Nan Zhao. 2013. A high-level energy consumption model for heterogeneous data centers. *Simulation Modelling Practice and Theory* 39 (2013).