profiles that are based on known specifications (e.g., specifications provided by a manufacturer of standardized computer hardware), but also device profiles based on virtual specifications. For example, through the use of device profiles based on virtual specifications, the systems and methods generate a "digital twin" of devices having non-homogeneous or unknown specifications. By using this digital twin, the system may complete incomplete sets of training data and thus properly train the artificial intelligence model.

[008] To generate the digital twin, the systems and methods may generate a network mapping of a computer network. The network mapping may indicate relationships between a plurality of devices in the computer network through the use of a series of interconnected nodes in which the nodes represent specific devices and the edges between nodes indicate the relationships. Furthermore, the network mapping may be populated with respective known specifications, wherein the respective known specifications are based on respective vendor information, respective model information, respective serial number information, or respective geographical location information. The system may then infer respective virtual specifications for any devices with non-homogeneous or unknown specifications based on the relationships between the plurality of devices and the respective known specifications. By doing so, the system may generate the needed virtual specifications that "fill in the gaps" of the network mapping and allow for an artificial intelligence model to be properly trained to provide power consumption predictions at an application-specific level even within network arrangements that feature devices with non-homogeneous or unknown specifications.

[008a] In an aspect, the present disclosure refers to a system for providing power consumption predictions for selected applications within network arrangements featuring devices with non-homogenous or unknown specifications, the system comprising:

one or more processors; and

a non-transitory, computer-readable medium comprising instructions recorded thereon that when executed by the one or more processors cause operations comprising:

generating a network mapping for a first network, wherein the network mapping indicates relationships between a plurality of devices in the first network, and wherein each of the plurality of devices is represented by respective nodes in the network mapping, wherein the

respective nodes comprise a first node subset and a second node subset, wherein the first node subset comprises nodes with respective known specifications, wherein the respective known specifications are based on respective vendor information, respective model information, respective serial number information, or respective geographical location information, wherein the second node subset comprises nodes with respective virtual specifications, and wherein the respective virtual specifications are inferred based on the relationships between the plurality of devices and the respective known specifications;

populating a device profile repository based on the network mapping;

receiving a first user request to generate a first application-specific power consumption prediction for a first application;

determining a first usage metric for the first application, wherein the first usage metric describes metadata related to one or more processes performed by the first application across a first network, wherein the first network comprises a plurality of devices;

determining, from the plurality of devices, a first subset of devices for serving the first application based on the first usage metric;

retrieving, from the device profile repository, a first subset of device profiles corresponding to devices in the first subset of devices;

determining, using a first artificial intelligence model, a first output, wherein the first artificial intelligence model is trained to generate application-specific power consumption predictions for respective applications based on inputted usage metrics for the respective applications and respective device profiles, from the device profile repository, for devices serving the respective applications; and

generating for display, on a user interface, the first application-specific power consumption prediction based on the first output.

[008b] In an aspect, the present disclosure refers to a method for providing power consumption predictions for selected applications within network arrangements featuring devices with non-homogenous or unknown specifications, the method comprising:

receiving a first user request to generate a first application-specific power consumption prediction for a first application;

determining a first usage metric for the first application, wherein the first usage metric describes metadata related to one or more processes performed by the first application across a first network, wherein the first network comprises a plurality of devices;

determining, from the plurality of devices, a first subset of devices for serving the first application based on the first usage metric;

retrieving, from a device profile repository, a first subset of device profiles corresponding to devices in the first subset of devices, wherein the device profile repository comprises a first set of device profiles corresponding to devices with known specifications and a second set of device profiles corresponding to devices with virtual specifications;

determining, using a first model, a first output, wherein the first model is trained to generate application-specific power consumption predictions for respective applications based on inputted usage metrics for the respective applications and respective device profiles, from the device profile repository, for devices serving the respective applications; and

generating for display, on a user interface, the first application-specific power consumption prediction based on the first output.

[009] In some examples, systems and methods for providing power consumption predictions for selected applications within network arrangements featuring devices with non-homogeneous or unknown specifications are described. For example, the system may receive a first user request to generate a first application-specific power consumption prediction for a first application. The system may determine a first usage metric for the first application, wherein the first usage metric describes metadata related to one or more processes performed by the first application across a first network, wherein the first network comprises a plurality of devices. The system may determine, from the plurality of devices, a first subset of devices for serving the first application based on the first usage metric. The system may retrieve, from a device profile repository, a first subset of device profiles corresponding to devices in the first subset of devices, wherein the device profile repository comprises a first set of device profiles corresponding to devices with known

WHAT IS CLAIMED IS:

1. A system for providing power consumption predictions for selected applications within network arrangements featuring devices with non-homogeneous or unknown specifications, the system comprising:

one or more processors; and

a non-transitory, computer-readable medium comprising instructions recorded thereon that when executed by the one or more processors cause operations comprising:

generating a network mapping for a first network, wherein the network mapping indicates relationships between a plurality of devices in the first network, and wherein each of the plurality of devices is represented by respective nodes in the network mapping, wherein the respective nodes comprise a first node subset and a second node subset, wherein the first node subset comprises nodes with respective known specifications, wherein the respective known specifications are based on respective vendor information, respective model information, respective serial number information, or respective geographical location information, wherein the second node subset comprises nodes with respective virtual specifications, and wherein the respective virtual specifications are inferred based on the relationships between the plurality of devices and the respective known specifications;

populating a device profile repository based on the network mapping;

receiving a first user request to generate a first application-specific power consumption prediction for a first application;

determining a first usage metric for the first application, wherein the first usage metric describes metadata related to one or more processes performed by the first application across a first network, wherein the first network comprises a plurality of devices;

determining, from the plurality of devices, a first subset of devices for serving the first application based on the first usage metric;

retrieving, from the device profile repository, a first subset of device profiles corresponding to devices in the first subset of devices;

determining, using a first artificial intelligence model, a first output, wherein the first artificial intelligence model is trained to generate application-specific power consumption predictions for respective applications based on inputted usage metrics for the respective

applications and respective device profiles, from the device profile repository, for devices serving the respective applications; and

generating for display, on a user interface, the first application-specific power consumption prediction based on the first output.

2. A method for providing power consumption predictions for selected applications within network arrangements featuring devices with non-homogeneous or unknown specifications, the method comprising:

receiving a first user request to generate a first application-specific power consumption prediction for a first application;

determining a first usage metric for the first application, wherein the first usage metric describes metadata related to one or more processes performed by the first application across a first network, wherein the first network comprises a plurality of devices;

determining, from the plurality of devices, a first subset of devices for serving the first application based on the first usage metric;

retrieving, from a device profile repository, a first subset of device profiles corresponding to devices in the first subset of devices, wherein the device profile repository comprises a first set of device profiles corresponding to devices with known specifications and a second set of device profiles corresponding to devices with virtual specifications;

determining, using a first model, a first output, wherein the first model is trained to generate application-specific power consumption predictions for respective applications based on inputted usage metrics for the respective applications and respective device profiles, from the device profile repository, for devices serving the respective applications; and

generating for display, on a user interface, the first application-specific power consumption prediction based on the first output.

3. The method of claim 2, wherein generating the device profile repository further comprises:

retrieving an aggregate carbon impact for a first device and a second device;

retrieving a first device profile for the first device, wherein the first device profile is based on a first known specification;

determining a first component carbon impact for the first device based on the first known

specification;

determining a difference between the first component carbon impact and the aggregate carbon impact;

determining a second component carbon impact for the second device based on the difference; and

determining a second virtual specification for the second device based on the second component carbon impact.

4. The method of claim 2, wherein generating the device profile repository further comprises:

retrieving an aggregate processing metric for a first device and a second device;

retrieving a first device profile for the first device, wherein the first device profile is based on a first known specification;

determining a first component processing metric for the first device based on the first known specification;

determining a difference between the first component processing metric and the aggregate processing metric;

determining a second processing metric for the second device based on the difference; and

determining a second virtual specification for the second device based on the second processing metric.

5. The method of claim 2, wherein generating the device profile repository further comprises:

retrieving a network mapping for the first network, wherein the network mapping indicates relationships between the plurality of devices, and wherein each of the plurality of devices is represented by respective nodes in the network mapping;

identifying a first node subset of the respective nodes, wherein the first node subset comprises nodes with respective known specifications;

generating a known network mapping based on populating the network mapping at each node of the first node subset with the respective known specifications;

identifying a second node subset of the respective nodes, wherein the second node subset comprises nodes with respective unknown specifications, and wherein the second node subset comprises unpopulated nodes in the known network mapping; and

inferring respective virtual specifications for each node in the second node subset based on the relationships between the plurality of devices and the respective known specifications in the known network mapping.

6. The method of claim 2, wherein determining the first usage metric for the first application further comprises:

determining a first process performed by the first application; and

determining a data volume created on a data array for consumption by the first application during the first process.

7. The method of claim 2, wherein determining the first usage metric for the first application further comprises:

determining a first process performed by the first application; and

determining watt-hour power consumption by the first application during the first process.

8. The method of claim 2, wherein determining the first usage metric for the first application further comprises:

determining a first process performed by the first application; and

determining a number of ports for switches dedicated to the first application during the first process.

9. The method of claim 2, wherein determining the first usage metric for the first application further comprises:

determining a first process performed by the first application; and

determining a number of data arrays created for the first application during the first process.

10. The method of claim 2, wherein determining the first subset of devices for serving the first application based on the first usage metric further comprises:

determining a number of devices of the plurality of devices required by the first application based on the first usage metric; and

determining the first subset of devices based on the number of devices.

11. The method of claim 2, wherein determining the first subset of devices for serving the first application based on the first usage metric further comprises:

determining a first process performed by the first application based on the first usage metric; and

filtering the plurality of devices based on whether a respective device is utilized during the first process.

12. The method of claim 2, wherein determining the first output further comprises:

determining a carbon impact for serving the first application with the first subset; and

determining the first application-specific power consumption prediction based on the carbon impact.

13. The method of claim 12, wherein determining the first application-specific power consumption prediction based on the carbon impact further comprises:

determining a power source for powering the first subset;

determining carbon impact ratio for the power source; and

adjusting the carbon impact based on the carbon impact ratio.

14. The method of claim 13, wherein determining the power source for powering the first subset further comprises:

determining a geographic location for the first subset; and

determining the power source based on the geographic location.

15. The method of claim 13, wherein determining the power source for powering the first subset further comprises:

determining a network location for the first subset; and

determining the power source based on the network location.