

SYSTEM AND METHOD FOR SIMULATING NETWORK EVENTS

BACKGROUND

5 The present invention relates to network simulation methods and systems.

Commercial-scale computer networks can be large and complex, and accordingly, it can be difficult to understand their behaviour using analytical methods such as those based on graph theory. For this reason, engineers or
10 technicians wishing to gain insight into network behaviour have more recently turned to network simulation software.

A particular problem which arises in computer networks used in the financial industry, such as card payment networks or Automated Clearing House (ACH)
15 systems, is that in addition to their complex topology, the nature of the data being transmitted across the network is highly confidential. Testing of algorithms (such as fraud detection algorithms) or network performance is therefore challenging, because of data security and privacy requirements. In addition, there can be data ownership issues since network participants may own the data
20 rather than it being centrally available. Payments data is extremely sensitive in nature and thus access to it, and subsequent analysis, requires strict compliance with the law, typically meaning that access can only be obtained in restricted physical or digital environments.

25 The present invention seeks to address the above difficulties.

SUMMARY

In one aspect of the present disclosure, there is provided a network simulation
30 system, comprising:

a request processor configured to receive input data relating to a topology of the network, the input data comprising node data indicative of respective resource levels of a plurality of nodes of the network, and relationship data indicative of respective relationship parameters of a

plurality of relationships between at least some of the nodes, the relationship parameters including at least a rate parameter and a transfer value parameter;

5 a simulation engine which is configured to output a stream of event data based on the relationship data, wherein respective events are generated by:

determining a source node and a destination node, and corresponding relationship parameters for the relationship between the source node and the destination node;

10 determining an event time based on the rate parameter of the corresponding relationship parameters;

determining a transfer value; and

responsive to a determination that the transfer value does not exceed a current resource level of the source node, 15 decrementing the current resource level of the source node by the transfer value, incrementing a current resource level of the destination node by the transfer value, and generating, at the event time, event data comprising a timestamp, a source node identifier of the source node, a destination node identifier of the 20 destination node, and the transfer value.

The simulation engine may be configured to spawn an event generation process for each of the plurality of relationships, and respective event generation processes may be configured to generate events for respective relationships.

25

In some embodiments, the request processor is configured to stream, to a remotely located user device, the event data.

The relationship parameters may include a transfer value parameter, and the 30 simulation engine may be configured to determine the transfer value based on the transfer value parameter.

In some embodiments, the transfer value parameter is a parameter of a transfer value distribution, and the simulation engine is configured to determine the transfer value by random sampling from the transfer value distribution.

- 5 The request processor may be configured to receive data relating to user-defined modifications to the network topology, and the simulation engine may be configured to send update data to one or more of the event generation processes, and/or spawn one or more additional event generation processes, according to the user-defined modifications. Accordingly, the user-defined
10 modifications can result in modification of the event data which is generated by the simulation engine.

- In certain embodiments, the simulation engine is configured to add or remove nodes, and/or add or remove relationships between nodes, and/or modify one or
15 more of the relationship parameters.

- The simulation engine, and/or individual event generation process(es), may be configured to sample the event time from a Poisson distribution having the rate parameter.
20

In some embodiments, the simulation engine is configured to output the event data in a standard format. The standard format may be ISO 20022 format.

- In some embodiments, at least one rate parameter is time-dependent.
25 Advantageously, this allows simulation of increased or decreased network traffic according to known fluctuations, thus allowing determination of the impact on network performance.

- In another aspect, the present disclosure relates to a network simulation
30 process, comprising:

receiving input data relating to a topology of the network, the input data comprising node data indicative of respective resource levels of a plurality of nodes of the network, and relationship data indicative of

respective relationship parameters of a plurality of relationships between at least some of the nodes, the relationship parameters including at least a rate parameter;

5 generating a stream of event data based on the relationship data, wherein respective events are generated by:

determining a source node and a destination node, and corresponding relationship parameters for the relationship between the source node and the destination node;

10 determining an event time based on the rate parameter of the corresponding relationship parameters;

determining a transfer value; and

15 responsive to a determination that the transfer value does not exceed a current resource level of the source node, decrementing the current resource level of the source node by the transfer value, incrementing a current resource level of the destination node by the transfer value, and generating, at the event time, event data comprising a timestamp, a source node identifier of the source node, a destination node identifier of the destination node, and the transfer value.

20

The process may comprise spawning an event generation process for each of the plurality of relationships, wherein respective event generation processes are configured to generate events for respective relationships.

25 The process may comprise streaming, to a remotely located user device, the event data.

The relationship parameters may include a transfer value parameter, and the transfer value may be determined based on the transfer value parameter.

30

The transfer value parameter may be a parameter of a transfer value distribution, and the transfer value may be determined by random sampling from the transfer value distribution.

The process may further comprise receiving data relating to user-defined modifications to the network topology, and sending update data to one or more of the event generation processes, and/or spawning one or more additional event generation processes, according to the user-defined modifications. The data relating to user-defined modifications may comprise data indicative of addition or removal of nodes, and/or addition or removal of relationships between nodes, and/or modification of one or more of the relationship parameters.

10 The event time may be sampled from a Poisson distribution having the rate parameter.

The process may comprise formatting the event data in a standard format, such as ISO 20022 format or another format suitable for financial network transaction messages. Advantageously, in conjunction with the other elements of the simulation process, it is thereby possible to generate realistic transaction data which can be used for downstream analysis and/or network performance testing without needing to deal with data privacy issues.

20 At least one rate parameter may be time-dependent.

In a further aspect, the present disclosure relates to a non-transitory computer-readable medium having instructions stored thereon which, when executed, cause at least one processor to perform a process according to any of the preceding paragraphs.

DETAILED DESCRIPTION

30 Embodiments of the invention will now be described with reference to the accompanying drawings, in which:

Figure 1 depicts an embodiment of a network simulation system in communication with a remotely located user;

Figure 2 is an exemplary architecture of a network simulation system;

Figure 3 is an example of a network topology suitable for simulation by the
5 network simulation system;

Figure 4 depicts the flow of an example network event generation process; and

Figure 5 is a flow diagram of an example network simulation process.

Embodiments generally relate to a network simulation system and process that
10 generate events, such as transactions, between nodes that are related by a computer network, such as a payments network. The events may be generated stochastically and/or in accordance with a fixed schedule. The event generation processes are balance-modifying, in that network events can result in changes to the available resources at respective nodes (such as available funds, in the
15 context of nodes that are accounts in a payments network). Additionally, event rates need not follow real-world time, and in particular can be generated at a much higher rate than real-world time.

Advantageously, the network topology can be modified by a user located
20 remotely from the network simulation system, for example via a RESTful API while an instance of the simulation is running, allowing complex, coordinated behaviours to be simulated, and to allow for locale-specific simulation.

Referring initially to Figure 1, there is shown in schematic form the architecture
25 of a network simulation system 100 which is in communication with a user device 102. The user device 102 may make requests to network simulation system 100, for example to start an instance of a network simulation, and/or to modify the network topology. The requests are received by the request processor 110 of network simulation system 100. Request processor 110 may
30 then invoke simulation engine 120 in order to generate network event data, which is subsequently returned to request processor 110 for output as a stream of events to user device 102. The rate of events and format of the event data may be controlled by request processor 110 and simulation engine 120 and may

be modified in accordance with requests from the user device 102 in a manner which will be described in more detail below.

Although the simulation system 100 is shown in communication with a remotely located user 102 in Figure 1, it is also possible for instances of the simulation to be started directly by a user at simulation system 100, or in automated fashion by a script executing on the simulation system 100, for example.

The architecture of the network simulation system 100 is shown in more detail in Figure 2. The network simulation system 100 may comprise at least one server 200 having at least one processor 202, volatile storage (RAM) 204, at least one network interface 206, and non-volatile storage (e.g., one or more magnetic disks and/or one or more solid-state drives) 220, all interconnected by a bus 210. A single server 200 is shown in Figure 2, though it will be appreciated that a network simulation system 100 may include multiple such servers.

The non-volatile storage 220 has stored thereon a plurality of modules, including operating system 222, request processor 110, and simulation engine 120 (and its sub-modules 122, 124, 126). Typically, these modules are implemented as software components, but it will be appreciated that they may encompass hardware, software, or a combination of hardware and software. As such, one or more of the modules 110, 120, 122, 124, 126 may exist as separate hardware components on the server 200 or on another server which is in communication with server 200, for example. Any of the modules shown in Figure 2 may be implemented in whole or in part as one or more application-specific integrated circuits (ASICs), field-programmable gate array (FPGAs), or other integrated circuit device in which the logic embodied in the module is implemented. In addition, the boundaries between modules as depicted are exemplary only, and it would be possible to implement alternative modules which perform the functions of any two or more of the modules shown in Figures 1 and 2.

Input data to the request processor 110 and/or the simulation engine 120 may optionally be stored in a data store 240 which resides on, or is in communication

with, storage medium 220. In particular, data store 240 may store data relating to the network topology (both as originally input and as modified during the course of a simulation), and to events generated by the simulation engine 120.

- 5 In order to generate network events, the request processor 110 is configured to receive (e.g., from user device 102) input data relating to a topology of the network. The input data comprises node data indicative of respective resource levels of a plurality of nodes of the network, and relationship data indicative of respective relationship parameters of a plurality of relationships between at least
- 10 some of the nodes. The relationship parameters for a given relationship between two nodes include at least a rate parameter, indicating a rate at which events between the two nodes occur, and a transfer value parameter, based on which the simulation engine 120 can determine the amount of a resource to transfer between the nodes. For example, in the context of a payments network,
- 15 the transfer value parameter may be a mean of a transaction value distribution of a certain mathematical form (e.g. a uniform or normal distribution), or may comprise values of an empirical distribution. The distribution may be sampled from, by the simulation engine 120, in order to determine the level of funds to be transferred between the nodes (accounts). In cases where the sampled value is
- 20 negative, and the simulated network is a payment network, the sampled value may be replaced by zero, a small non-zero value (such as £0.01), or some other predefined minimum transaction value. Alternatively, the transfer value parameter may be a fixed value.
- 25 The simulation engine 120 may be configured to spawn an event generation process for each relationship defined by the network topology. For example, as shown in Fig. 1, the network topology comprises N relationships, each associated with a respective event generation process 130.1, 130.2, ..., 130.N. Each event generation process 130.1, ..., 130.N has as input the parameters of
- 30 the network relationship with which it is associated, and as output a stream of simulated network events, which are streamed to simulation engine 120 and then in turn to request processor 110, which may in turn stream events to user device 102.

An example of an input network topology is shown in Figure 3, in which three nodes A, B and C are shown. Nodes A and B have a relationship characterised by a parameter λ_{ab} , nodes B and C have a relationship characterised by λ_{bc} , and
5 nodes A and C have a relationship characterised by λ_{ca} . Each node A, B and C has a starting resource level, which can be provided via user input, for example. The respective λ s are rate parameters for the respective relationships. In general, each relationship will also have a transfer value parameter or parameters associated with it, as discussed above. For example, relationship A-
10 B may be associated with a single fixed transfer value, relationship B-C may be associated with an empirical distribution of transfer values that can be sampled from (e.g. a set of real transaction values associated with a bank account over a period of time, such as a year), and relationship C-A may be associated with transfer parameters (μ , σ) which are the mean and standard deviation of a
15 normal distribution of transfer values.

It will be appreciated that the example shown in Figure 3 is illustrative only, and that a realistic simulated network topology may contain hundreds, thousands or even tens of thousands of nodes and relationships.

20

The input data representing the network topology may be provided to the request processor 110 in various ways. For example, an input file containing the initial node data and relationship data may be read by the request processor 110 and the various parameters stored in data store 240. Alternatively, a user device
25 102 may transmit requests to the request processor 110 to generate the desired network structure. This may be done interactively by the user, or more typically, as a batch, for example using a script or other means of automation such as a wrapper function which has input parameters such as the number of nodes to be created, a list of nodes which are to have relationships, and so on.

30

In order to generate a network topology that provides as realistic an approximation as possible to the network being simulated, the network simulation system 100 (for example, via request processor 110), or a module

executing on the user device 102, may employ a random graph generation algorithm which is configured to reproduce the desired network characteristics. For example, the algorithm may generate graphs according to the Watts-Strogatz model, the Erdős–Rényi model, or the Barabási–Albert model. In some
5 embodiments, the network simulation system 100 may receive empirical network topology data relating to one or more known networks of the appropriate type (e.g. payment networks), and derive, from the empirical network topology data, suitable parameters as input for the random graph generation algorithm.

10 In some embodiments, network relationships can be initialised in such a way that they are mutually independent of other relationships. Accordingly, the event generators spawned by simulation engine 120 produce output streams that are also independent of each other.

15 In other embodiments, it may be desirable to generate event data in an order which is dictated by the network hierarchy, such that there are dependencies between at least some of the event generation processes. For example, in the context of simulated fraud in a payments network, the network hierarchy may comprise a first account (node) associated with a victim being defrauded, and a
20 plurality of other accounts (nodes) to which the victim's funds are dispersed.

In order to implement conditionality for a network relationship, simulation engine 120 may define a conditionality parameter as one of the relationship parameters. For example, the conditionality parameter c may lie in the range from 0 to 1, and
25 the associated event generation process 130.1 may sit in a "wait" mode until a proportion c of inbound relationships (i.e., relationships for which the target node has the same identifier as the source node of the relationship in question) generates transaction events before its own transaction generation will begin. This mechanism allows for fully or partial conditional networks to be
30 implemented alongside stochastic relationships. The conditional relationships could also change through the course of the simulation, by updating the conditional parameter, thereby effectively changing the way funds were flowing around the network. The total number of inbound relationships, and the number

which have transacted, may be passed by simulation engine 120 as input parameters to the event generation process 130.1.

5 In other embodiments, event generation process 130.1 may sit in “wait” mode until a specific inbound relationship, or set of inbound relationships, has generated one or more transaction events.

10 In certain embodiments, such as those relating to payment networks, it may be desirable to associate additional data with nodes and/or relationships in the network. For example, if each node is a bank account then it may be desirable to associate a person’s name (such as “John Smith”) with each node, as well as a bank name (such as “RBS”). In this case, in order to generate more realistic node data, the additional data may be sampled from appropriate empirical distributions such as census-derived name distributions and known distributions
15 of bank accounts by market share. Other types of data associable with nodes or relationships, for which empirical distributions may be used to generate suitable values, include industry codes, sort codes (distributed by geography and/or bank code), bank balances, and activity level (for example, an account used to perpetrate fraud will typically have a much higher activity level than a normal
20 account).

The request processor 110 may be configured to receive requests to modify the network topology, for example from a remotely located user 102 by way of a RESTful API. For example, user 102 may add a node (account) to the network
25 by a command such as:

```
curl <server_name>:port/account -d {"SortCode":"00-00-00",  
"AccountNumber":"12345678","Name":"Mr John Smith","Balance":"9999"}
```

30 This creates an account having a sort code of 00-00-00, an account number of 12345678, and a balance of 9999 in the currency of the simulation (for example, GBP), and having an account name of Mr John Smith. Here, <server_name> is the server on which the instance of the simulation is being run, and <port> is the

port on that server over which the user device 102 is connected or authorised to connect. In some embodiments, a bank code may also be associated with the created node, though it will be appreciated that a sort code may act as a proxy for a bank code.

5

In another example, user 102 may create a relationship between a source node having an ID of 00-00-0012345678 and a destination node having an ID of 11-11-1187654321 by a command such as the following:

```
10 curl <server_name>:port/relationship -d {"SourceId":"00-00-0012345678","TargetId":"11-11-1187654321","Rate":"9.99"}
```

Here, the source and destination node identifiers have been generated by concatenating the respective sort code and account number that were used to
15 create the respective nodes. The Rate specified in the above command is the rate parameter for the relationship (i.e., λ as discussed above).

Once created, nodes or relationships may be modified, via API calls analogous to the above, by using the respective node identifier or relationship identifier
20 (which may be a concatenation or other unique combination of the source node identifier and target node identifier, for example). For example, a node may be deleted via a HTTP DELETE call using the node identifier, or may be updated via a HTTP PUT call using the node identifier.

25 Thus, as described above, an instance of the network event simulator can be run on a server 100 and users 102 can then interact with the underlying simulation as it is running. The simulation functions may allow for accounts and relationships to be added, removed or modified via API calls. This advantageously allows users 102 to change both account and relationship level
30 properties of the underlying network in real time.

The simulation engine 120 is configured to produce a stream of network events, such as transactions, for each relationship in the simulated network. In certain

embodiments, each relationship (and thus each event generation process 130.1, ..., 130.N) is associated with an underlying point process, such as a Poisson process, and the time between transactions is calculated depending on the underlying point process. After the correct amount of time according to the simulation has passed, a transaction value is generated and a request is made as to whether the transaction can occur. This is governed according to the source node balance (resource level); if it has sufficient funds then the transaction is allowed to occur, otherwise it is rejected.

10 An exemplary network simulation process 500 is shown in Figures 4 and 5. In Figure 4, a schematic representation of the process flow is shown, in which nodes A and B have a relationship characterised by a rate parameter λ_{AB} .

Referring to Figure 5, the process 500 begins at step 502 with the system 100 receiving input data relating to a topology of the network. This may be from user device 102, via request processor 110, for example. The input data comprises node data indicative of respective resource levels, such as account balances, of a plurality of nodes of the network, and relationship data indicative of respective relationship parameters of a plurality of relationships between at least some of the nodes, the relationship parameters including at least a rate parameter and a transfer value parameter.

At step 504, the process determines a source node, such as node A, and a destination node, such as node B, and corresponding relationship parameters for the relationship between the source node A and the destination node B. The relationship parameters may include the rate parameter λ_{AB} and a parameter of a distribution of transfer values, for example (e.g., μ and σ of a normal distribution as discussed above in relation to Figure 3). At this point, the process may spawn an event generation process, such as event generation process 130.1, for the AB relationship.

At step 506, the event generation process 130.1 may invoke the sampling module 124 of simulation engine 120 to sample a transaction time t (more

specifically, a time to the next transaction) from the underlying point process associated with the AB relationship, for example a Poisson process with rate parameter λ_{AB} .

- 5 Next, at step 508, the sampling module 124 samples a transaction value v using the transfer value parameter(s) of the AB relationship.

At step 510, the event generation process 130.1 requests, at transaction module 122, a transaction at time t with value v from the source node A to the
 10 destination node B. Transaction module 122 checks the transfer value v against the resource level (balance) of node A.

If the resource level is sufficient, then the transaction is accepted, and at step 512, a transaction message is generated. The transaction message may include
 15 at least a timestamp, a source node identifier of the source node, a destination node identifier of the destination node, and the transfer value v . The transaction message may also include an indication that the transaction event is a synthetic event (i.e., not generated by a real payment network), and may include
 20 versioning information to indicate which version of the simulator was used to generate it. The transaction module 122 also updates the resource levels of nodes A and B by decrementing node A by v and incrementing node B by v .

At step 514, the transaction data may be passed to the request processor 110 to
 25 be transmitted to user device 102.

As mentioned above, a Poisson point process can be used to model the inter-arrival time between transactions and is described by a single parameter λ . For a relationship represented as a Poisson process with parameter λ in a simulation started at time $t = 0$, each subsequent inter-arrival time is generated
 30 by sampling exponential random numbers with parameter λ . In practice this can be related to relationships that transact with a relatively fixed frequency i.e. if a source account pays a destination account X times per time period, such as in a month, then setting $\lambda = X$ will result in a stream of transactions that, as a long-

term average, will produce X number of transactions between source and destination for every month in the simulation.

In some embodiments, the underlying process associated with a relationship may be non-Poisson, or may be a sum of different types of processes (including Poisson and non-Poisson). For example, the underlying process may be a sum of a Poisson process, representing randomly distributed transaction times, and a process which produces events at fixed time intervals (e.g., once per month, indicating a regular monthly payment). Accordingly, it is possible to include various types of fixed frequency transactions e.g. transactions similar to direct debits and standing orders in nature. In combination this may enable all users to build rich dynamical datasets that can simulate known behaviours.

The sum of Poisson processes with parameters $\lambda_1, \lambda_2, \lambda_3, \dots, \lambda_n$ is itself a Poisson process with

$$\lambda = \sum_{i=1}^n \lambda_i$$

Accordingly, it is possible to tune a simulated dataset to a certain desired volume of transactions over timescales of interest in a very elegant way. In particular, for use cases such as determining the ability of a network to cope with a certain transaction volume (both in terms of bandwidth and storage), where rich dynamics are not required, a network with n relationships obeying the above equation is sufficient.

The above can be modified to include dynamics such that λ is no longer constant or stationary, but rather varies with time i.e. $\lambda \rightarrow \lambda(t)$. Advantageously, the form of $\lambda(t)$ can then be modelled to capture variations in the transaction rate. For example around the end of the tax year or at Christmas and New Year the rate of transactions might increase or decrease depending on the payment network being considered.

Advantageously, in the context of payment networks, the simulation system is not constrained to replicate any particular payment system, since it generates transaction events based on collections of point processes. In this way the simulator is completely agnostic to the type of dataset being constructed.

- 5 Optionally, the generated transactions may be transformed to be compliant with a particular payment data standard, such as ISO8583 or ISO20022. This can be achieved via a data standards module 126 of the simulation engine 120. By making the output compliant with a standard such as ISO8583 or ISO20022, the user 102 has the flexibility to simulate any payment system they wish, adapting
- 10 the simulation output to their particular needs. A data element of the ISO8583 or ISO20022 message may carry the indication that the event is synthetic and/or the versioning information for the simulator.

Embodiments may have one or more of the following advantages:

15

1. Enabling testing of algorithms. One of the key advantages of a synthetic dataset generated according to the process described above is that it enables testing of algorithms, such as fraud detection algorithms, on realistic data. For example, the network simulation system allows modification of the network
- 20 topology such that particular nodes are much more active than usual, a known indicator of fraud in the context of payment networks. Accordingly, a stream of transactions incorporating such known fraudulent accounts can be passed to a fraud detection algorithm in order to test whether the algorithm is capable of detecting this injected signal. In addition, due to the flexibility of the framework
- 25 described above, various scenarios can be simulated in order to help refine and improve existing algorithms.

2. Facilitating developer testing. Since the system and process described above enable generation of realistic data sets on short time scales, and may be made
- 30 e.g. ISO20022 compliant without being restricted to a particular payments framework, users 102 may generate sufficient data volumes adapted to their desired format in order to perform stress testing of a network. Such stress testing may include tests such as checking whether messages are of the correct

format, whether algorithms designed to check network integrity find mistakes in network messages, whether peak day volumes affect the performance of software, and whether there is sufficient database storage to cope with surges in network traffic.

5

3. Working with Sensitive Data. To be able to work with a dataset that is synthetic circumvents the need to be in a secure environment to build and test algorithms. Much of the testing work can be done using the synthetic data as input, and once the initial testing has shown that the algorithms are working as expected, the code can be tested against real data in the secure environment.

10

4. Visualisation & Demos. Having synthetic data allows generation of very rich and detailed visualisations and demos that can be used to demonstrate algorithm results.

15

5. Proofs of Concept. As well as testing algorithms as mentioned above, having a synthetic dataset allows for more exploratory work to be done without the need for permissions to be sought first.

20 Ultimately, embodiments of the invention enable a secure, principled approach to simulating network data which gives users the freedom to interface with the simulator science focused on payments or other sensitive network data that would otherwise present enormous privacy and security risks.

25

CLAIMS

1. A system for simulating network events, comprising:
 - a user device;
 - 5 a request processor of a network simulation system operably connected to the user device, wherein the request processor is configured to receive a request for network simulation data and to receive input data from the user device, the input data relating to user-defined modifications to a topology of a network, the input data comprising:
 - 10 node data indicative of respective resource levels of a plurality of nodes of the network, and
 - relationship data indicative of respective relationship parameters of a plurality of relationships between at least some of the nodes, the relationship parameters including a rate parameter and a transfer value
 - 15 parameter; and
 - a simulation engine of the network simulation system operably connected to the request processor, wherein the simulation engine is configured to output a stream of event data to the request processor, the event data based on the relationship data, and wherein the simulation engine
 - 20 generates respective events of the stream of event data by:
 - determining a source node and a destination node, and corresponding relationship parameters for relationship between the source node and the destination node;
 - determining an event time based on the rate parameter of the
 - 25 corresponding relationship parameters;
 - determining a transfer value based on the transfer value parameter and a random sampling from a transfer value distribution;
 - responsive to a determination that the transfer value does not exceed a current resource level of the source node,
 - 30 decrementing the current resource level of the source node by the transfer value,
 - incrementing a current resource level of the destination node by the transfer value, and

generating, at the event time, event data comprising a timestamp, a source node identifier of the source node, a destination node identifier of the destination node, and the transfer value;
spawning an event generation process for each of the plurality of relationships, wherein respective event generation processes are configured to generate events for respective relationships;
transmitting update data to at least some of the event generation processes based on the user-defined modifications;
spawning at least one additional event generation process based on the update data; and
streaming, by the request processor, the event data to the user device.

2. The network simulation system according to claim 1, wherein the simulation engine is configured to add or remove nodes, and/or add or remove relationships between nodes, and/or modify one or more of the relationship parameters.

3. The network simulation system according to claim 1, wherein the simulation engine is configured to sample the event time from a Poisson distribution having the rate parameter.

4. The network simulation system according to claim 1, wherein the rate parameter is time-dependent.

5. A network event simulation process, comprising:
receiving, by a network simulation system comprising a request processor and a simulation engine, a request for network simulation data and input data from a user device, the input data relating to user-defined modifications to a topology of a network, the input data comprising node data indicative of respective resource levels of a plurality of nodes of the network, and relationship data indicative of respective relationship parameters of a plurality of relationships between at least some of the nodes, the relationship parameters including a rate parameter and a transfer value parameter;

generating, by the simulation engine of the network simulation system, a stream of event data based on the relationship data, wherein respective events of the stream of event data are generated by:

5 determining a source node and a destination node, and corresponding relationship parameters for relationship between the source node and the destination node;

determining an event time based on the rate parameter of the corresponding relationship parameters;

10 determining a transfer value based on the transfer value parameter and a random sampling from a transfer value distribution;

responsive to a determination that the transfer value does not exceed a current resource level of the source node,

decrementing the current resource level of the source node by the transfer value,

15 incrementing a current resource level of the destination node by the transfer value,

generating, at the event time, event data comprising a timestamp, a source node identifier of the source node, a destination node identifier of the destination node, and the transfer value;

20 spawning an event generation process for each of the plurality of relationships, wherein respective event generation processes are configured to generate events for respective relationships;

transmitting update data to at least some of the event generation processes based on the user-defined modifications;

25 spawning at least one additional event generation process based on the update data; and

streaming, by the request processor of the network simulation system, the event data to the user device.

30 6. The network simulation process according to claim 5, wherein the data relating to user-defined modifications comprises data indicating addition or removal of nodes, and/or addition or removal of relationships between nodes, and/or modification of one or more of the relationship parameters.

7. The network simulation process according to claim 5, wherein the event time is sampled from a Poisson distribution having the rate parameter.
- 5 8. The network simulation process according to claim 5, wherein the rate parameter is time-dependent.
9. A non-transitory computer-readable medium having instructions stored thereon which, when executed, cause at least one processor to:
 - 10 receive a request for network simulation data and input data from a user device relating to user-defined modifications to a topology of a network, the input data comprising node data indicative of respective resource levels of a plurality of nodes of the network, and relationship data indicative of respective relationship parameters of a plurality of relationships between at least some of the nodes, the
 - 15 relationship parameters including a rate parameter and a transfer value parameter;
 - generate a stream of event data based on the relationship data, wherein respective events of the stream of event data are generated by:
 - determining a source node and a destination node, and
 - 20 corresponding relationship parameters for relationship between the source node and the destination node;
 - determining an event time based on the rate parameter of the corresponding relationship parameters;
 - determining a transfer value based on the transfer value parameter
 - 25 and a random sampling from a transfer value distribution;
 - responsive to a determination that the transfer value does not exceed a current resource level of the source node,
 - decrementing the current resource level of the source node by the transfer value,
 - 30 incrementing a current resource level of the destination node by the transfer value, and
 - generating, at the event time, event data comprising a timestamp, a source node identifier of the source node, a

destination node identifier of the destination node, and the transfer value;

spawning an event generation process for each of the plurality of relationships, wherein respective event generation processes are configured to generate events for respective relationships;

5

transmitting update data to at least some of the event generation processes based on the user-defined modifications;

spawning at least one additional event generation process based on the update data; and

10

stream the event data to the user device.