



# SimPy Documentation

*Versão 3.0.10*

**Team SimPy**

03/06/2018

<b>1</b>	<b>Sumário</b>	<b>1</b>
<b>2</b>	<b>SimPy em 10 Minutos</b>	<b>3</b>
2.1	Instalação . . . . .	3
2.2	Conceitos Básicos . . . . .	4
2.3	Esperando por um Processo . . . . .	5
2.4	Interrompendo a execução de outro Processo . . . . .	6
2.5	E agora? . . . . .	7
2.6	Recursos Compartilhados . . . . .	7
2.7	Como proceder . . . . .	9
<b>3</b>	<b>Topical Guides</b>	<b>10</b>
3.1	Conceitos básicos do SimPy . . . . .	10
3.2	Environments . . . . .	11
3.3	Events . . . . .	14
3.4	Process Interaction . . . . .	18
3.5	Shared Resources . . . . .	21
3.6	Real-time simulations . . . . .	29
3.7	Monitorando . . . . .	30
3.8	Time and Scheduling . . . . .	35
3.9	Porting from SimPy 2 to 3 . . . . .	36
<b>4</b>	<b>Examples</b>	<b>41</b>
4.1	Condition events . . . . .	41
4.2	Interrupts . . . . .	41
4.3	Monitoring . . . . .	41
4.4	Resources: Container . . . . .	41
4.5	Resources: Preemptive Resource . . . . .	41
4.6	Resources: Resource . . . . .	41
4.7	Resources: Store . . . . .	42
4.8	Shared events . . . . .	42
4.9	Waiting for other processes . . . . .	42
4.10	All examples . . . . .	42
<b>5</b>	<b>API Reference</b>	<b>61</b>
5.1	<code>simpy</code> . . . . .	61
5.2	<code>simpy.core</code> — SimPy's core components . . . . .	62
5.3	<code>simpy.exceptions</code> — Exception types used by SimPy . . . . .	64

5.4	<code>simpy.events</code> — Core event types . . . . .	64
5.5	<code>simpy.resources</code> — Shared resource primitives . . . . .	72
5.6	<code>simpy.rt</code> — Real-time simulation . . . . .	79
5.7	<code>simpy.util</code> — Utility functions for SimPy . . . . .	80
<b>6</b>	<b>About SimPy</b>	<b>81</b>
6.1	SimPy History & Change Log . . . . .	81
6.2	Acknowledgments . . . . .	91
6.3	Ports and comparable libraries . . . . .	91
6.4	Defense of Design . . . . .	91
6.5	Release Process . . . . .	94
6.6	License . . . . .	97
<b>7</b>	<b>Indices and tables</b>	<b>98</b>
	<b>Índice de Módulos do Python</b>	<b>99</b>

---

## Sumário

---

Simpy é um framework baseado em processos com simulação de eventos discretos utilizando as bibliotecas padrão do Python.

Os processos em Simpy são definidos através de [funções generator](#) e podem, por exemplo, ser utilizadas para definir componentes como compradores, veículos ou agentes. Simpy também disponibiliza diversos tipos de [recursos compartilhados](#) para utilização de recursos propensos a congestionamento ( como servidores, caixas e túneis)

As simulações podem ser realizadas de modo “*mais rápido possível*”, em *tempo real* (de acordo com o desenrolar do tempo) ou manualmente *passo a passo* por cada evento.

Though it is theoretically possible to do continuous simulations with SimPy, it has no features that help you with that. On the other hand, SimPy is overkill for simulations with a fixed step size where your processes don't interact with each other or with shared resources.

A short example simulating two clocks ticking in different time intervals looks like this:

```
>>> import simpy
>>>
>>> def clock(env, name, tick):
...     while True:
...         print(name, env.now)
...         yield env.timeout(tick)
...
>>> env = simpy.Environment()
>>> env.process(clock(env, 'fast', 0.5))
<Process(clock) object at 0x...>
>>> env.process(clock(env, 'slow', 1))
<Process(clock) object at 0x...>
>>> env.run(until=2)
fast 0
slow 0
fast 0.5
slow 1
fast 1.0
fast 1.5
```

A documentação contém um [tutorial](#), [diversos guias](#) explicando os conceitos chave, uma boa gama de [exemplos](#) e a [documentação da API](#).

SimPy is released under the MIT License. Simulation model developers are encouraged to share their SimPy modeling techniques with the SimPy community. Please post a message to the [SimPy mailing list](#).

There is an introductory talk that explains SimPy's concepts and provides some examples: [watch the video](#) or [get the slides](#).

SimPy has also been reimplemented in other programming languages. See the *[list of ports](#)* for more details.

---

## SimPy em 10 Minutos

---

Nesta seção você aprenderá o básico sobre Simpy em apenas alguns minutos. Ao término, estarás apto a implementar uma simulação simples usando Simpy e você será capaz de decidir se o Simpy é o que necessita para seu projeto. Também lhe daremos algumas dicas de como implementar simulações mais complexas.

### Instalação

SimPy é implementado sem necessidade de dependências ou recursos além do próprio Python. Simpy roda em Python 2 ( $\geq 2.7$ ) e Python 3 ( $\geq 3.2$ ). PyPy também é suportado. Se já tiver instalado o `pip`, só precisa digitar

```
$ pip install simpy
```

e pronto.

### Instalando a partir do código fonte

Uma alternativa a instalação é fazer o [download do SimPy](#) e instalá-lo manualmente. Descompacte o arquivo, abra um terminal no diretório onde o Simpy foi descompactado e digite:

```
$ python setup.py install
```

Nesse modelo de instalação ainda é possível executar os testes do Simpy para verificar se tudo está funcionando de acordo. Para isto você precisa do `pytest`. Execute o comando a seguir diretamente do diretório onde está o código fonte do Simpy:

```
$ py.test --pyargs simpy
```

### Atualizando o Simpy 2 para versões mais recentes

Se você já está familiarizado com o Simpy 2, por favor leia o Guia *Porting from SimPy 2 to 3*.

### O que vem a seguir?

Após a instalação do Simpy, você provavelmente deseja simular algo. Na *próxima seção* você conhecerá os conceitos básicos do Simpy.

## Conceitos Básicos

Simpy é uma biblioteca de simulação de eventos discretos. O comportamento de componentes ( como **veículos**, **clientes** ou **mensagens**) é modelado como *processes*. Todos os *processes/processos* estão em *environment/ambiente*. Eles interagem com o ambiente e um com o outro através de *eventos*.

Os *processos* são definidos através de simples Python *generators*. Você pode chamar as *funções de processos* ou *métodos dos processos*, da mesma maneira que chama uma função normal ou um método de uma classe. Durante seu ciclo de vida, eles criam eventos e os executa de acordo com a ordem que deverão ser acionados.

Quando um processo tem o retorno de um evento, o processo fica em modo *suspended/suspense*. Simpy *resumes/continua* a execução do processo, quando o evento ocorre ( é dito que o evento é *triggered/acionado*). Diversos processos podem aguardar pelo mesmo evento. Simpy resume/aciona todos os eventos na mesma ordem em que foram retornados.

Eventos do tipo *Timeout*, são acionados após um determinado tempo ( da simulação ). Essa funcionalidade permite que um processo fique aguardando apenas por um tempo definido. Um *Timeout* e todos os outros eventos podem se criados com a execução de um método do tipo *Environment* no qual o processo se encontra (*Environment.timeout()* por exemplo).

### Primeiro Processo

Nosso primeiro exemplo será um processo *car*. O *car/carro* alternará entre ser dirigido(*drive*) e estar estacionado (*park*) por um tempo. Quando ele estiver em movimento ( ou estacionando ), será impresso o tempo atual da simulação. Ou seja, sempre que o *estado* do processo mudar.

Então vamos começar:

```
>>> def car(env):
...     while True:
...         print('Estacionando a %d' % env.now)
...         parking_duration = 5
...         yield env.timeout(parking_duration)
...
...         print('Dirigindo em %d' % env.now)
...         trip_duration = 2
...         yield env.timeout(trip_duration)
```

O processo *car* requer uma referencia para a classe *Environment* (*env*) de modo ordenado para a criação/definição dos novos eventos relacionados a ele. O comportamento do *car* é descrito em um loop infinto. Lembrando que está função é um *generator*. Apesar dele nunca terminar (loop infinito), ele entrega o controle do fluxo das ações sempre que encontra a declaração *yield*. Uma vez que o evento(*yield*) é finalmente acionado(triggered) ( ou ocorre ), a simulação continua a execução do procedimento.

Como dito antes, nosso carro/*car* alterna entre estar estacionando/*parking* e estar em movimento/*driving*. É anunciado seu novo status imprimindo uma mensagem e o tempo atual da simulação ( retornado pela propriedade *Environment.now*). Então é chamada a função *Environment.timeout()* para criar um evento do tipo *Timeout*. Este evento determina o ponto exato em que o carro deixou de estacionar/*parking* ou de estar em movimento/*driving*, respectivamente. Através do evento 'yielding', é informado para a simulação que será aguardada a mudança de estado.

Já que o comportamento do nosso carro foi definido/modelado, vamos criar uma instância dele e ver como ele se comporta:

```
>>> import simpy
>>> env = simpy.Environment()
>>> env.process(car(env))
```

```

<Process(car) object at 0x...>
>>> env.run(until=15)
Estacionado em 0
Dirigindo em 5
Estacionando em 7
Dirigindo em 12
Estacionando em 14

```

A primeira coisa que fazemos é criar uma instância da classe `Environment`. Essa instância ( no caso `env` ) é então passada para o processo `car` em `car(env)`. Que por sua vez cria um *process generator* que para ser iniciado tem de ser declarado na instância `env`, para isso se usa o método `Environment.process()`.

Observer, que neste momento, nada do código/procedimento definido em nosso processo está sendo executado ainda. Sua execução é agendada de acordo com o tempo da simulação

O `Process` retornado por `process()` pode ser usado para interagir com o processo ( na próxima seção veremos mais sobre isso).

E finalmente a simulação começa quando é executado `run()` e definimos quando será o término da simulação.

## E agora?

Já que você se familiarizou com a terminologia utilizada no SimPy e seus conceitos básicos. Na [próxima seção](#), vamos entender como funciona a interação do processo.

A instância de `Process` retornada por `Environment.process()` pode ser usada para interagir com o processo. `Environment.process()` can be utilized for process interactions. Dois exemplos mais comuns de uso são para aguardar o término de algum outro processo e interromper a execução de outro processo enquanto se espera por um evento.

## Esperando por um Processo

Como mencionado, um `Process` SimPy pode ser usado como um evento ( tecnicamente falando, um processo é um evento). Se você marcá-lo com `yield`, ele continua uma vez que o processo tenha concluído. Imagine a simulação de uma *lava a jato* onde os carros entram para a limpeza e aguardam que o processo de lavagem seja concluído. Ou uma simulação de um *aeroporto* onde os passageiros tem de esperar até que os procedimentos de segurança sejam concluídos.

Vamos assumir que o carro do nosso exemplo anterior, vire do nada um carro elétrico. Carros elétricos geralmente levam muito tempo carregando suas baterias após uma viagem. Eles teriam de esperar até que suas baterias estivessem carregadas para que continuassem sua viagem.

Poderíamos alterar o modelo adicionando ao carro o processo `charge()`. Antes, é preciso refatorar nossa classe carro para que os processos tenham dois métodos: `run()` ( o qual antes era a função `car()` ) e `charge()`.

O processo `run` é iniciado automaticamente quando `Car` é instanciado. E um processo `charge` é chamado toda vez que o carro começa a estacionar. Ao se solicitar uma instância de `Process` retornado pelo método `Environment.process()`, o processo `run` fica aguardando o término do evento:

```

>>> class Car(object):
...     def __init__(self, env):
...         self.env = env
...         # Start the run process everytime an instance is created.
...         self.action = env.process(self.run())
...
...     def run(self):

```

```

...     while True:
...         print('Estacionando e carregando em %d' % self.env.now)
...         charge_duration = 5
...         # We yield the process that process() returns
...         # to wait for it to finish
...         yield self.env.process(self.charge(charge_duration))
...
...         # Ao termino do carregamento então podemos
...         # voltar a dirigir.
...         print('Dirigindo em %d' % self.env.now)
...         trip_duration = 2
...         yield self.env.timeout(trip_duration)
...
...     def charge(self, duration):
...         yield self.env.timeout(duration)

```

Iniciando a simulação novamente: Nós criamos um ambiente/*environment*, um ( ou mais ) carros e finalmente chamamos o método `run()`.

```

>>> import simpy
>>> env = simpy.Environment()
>>> car = Car(env)
>>> env.run(until=15)
Estacionando e carregando em 0
Dirigindo em 5
Estacionando e carregando em 7
Dirigindo em 12
Estacionando e carregando em 14

```

## Interrompendo a execução de outro Processo

Digamos que você não quer aguardar até que o carro elétrico esteja completamente carregado e sim deseja parar o processo de carga para começar a dirigir.

SimPy permite que você interrompa a execução de um processo chamando o método `interrupt()`

```

>>> def driver(env, car):
...     yield env.timeout(3)
...     car.action.interrupt()

```

O processo `driver` tem uma referência para processo `action` do carro. Depois de esperar por 3 passos da simulação, ele interrompe aquele processo.

As interrupções são aplicadas nos processos como exceções `Interrupt` que podem (ou devem) ser tratadas pelo processo que foi interrompido. O processo por sua vez pode então decidir o que fazer em seguida (por exemplo: continuar a esperar o evento que o originou ou iniciar um novo evento):

```

>>> class Car(object):
...     def __init__(self, env):
...         self.env = env
...         self.action = env.process(self.run())
...
...     def run(self):
...         while True:
...             print('Estacionando e carregando em %d' % self.env.now)
...             charge_duration = 5
...             # Caso seja interrompido enquanto carrega a bateria

```

```

...         try:
...             yield self.env.process(self.charge(charge_duration))
...         except simpy.Interrupt:
...             # Ao receber a solicitacao de interrupcao do processo
...             # de carga, nos paramos de carregar e mudamos o estado
...             # para "dirigindo"/"driving"
...             print('Interrompido o processo. Espero que a bateria tenha carga suficiente ...')
...             print('Dirigindo em %d' % self.env.now)
...             trip_duration = 2
...             yield self.env.timeout(trip_duration)
...
...     def charge(self, duration):
...         yield self.env.timeout(duration)

```

Se compararmos as mensagens de saída da simulação anterior com a atual, veremos que o carro agora começa a ser dirigido em 3 ao invés de 5:

```

>>> env = simpy.Environment()
>>> car = Car(env)
>>> env.process(driver(env, car))
<Process(driver) object at 0x...>
>>> env.run(until=15)
Estacionando e carregando em 0
Interrompido o processo. Espero que a bateria tenha carga suficiente ...
Dirigindo em 3
Estacionando e carregando em 5
Dirigindo em 10
Estacionando e carregando em 12

```

## E agora?

Nós demonstramos apenas dois métodos simples para interação com os processos—esperando por um processo e interrompendo um processo. Dê uma olhada em [Topical Guides](#) ou documentação da API *Process* para maiores detalhes.

Na [próxima seção](#) vamos entender o básico da utilização de recursos compartilhados.

## Recursos Compartilhados

SimPy oferece 3 tipos de *resources* que o ajudam a modelar suas simulações, onde múltiplos processos podem usar um recurso com capacidade limitada ( ex., carros em um posto com uma limitada quantidade de bombas ) ou o clássico problema produtor-consumidor.

Nesta seção, faremos uma breve introdução da classe *Resource* do SimPy.

### Uso básico da classe Resource

Faremos uma modificação simples no nosso processo do carro elétrico `car` que foi criado na última seção.

O carro agora irá para um *estação de carga da bateria*(`bcs`) e solicitar 1 dos 2 pontos de carga ( *charging spots* ). Se todos os pontos estiverem em uso, ele aguarda a liberação de um deles. Só então, obviamente, deve começar a carregar a bateria e assim deixar a estação:

```

>>> def car(env, name, bcs, driving_time, charge_duration):
...     # Simular seguir para a BCS
...     yield env.timeout(driving_time)
...
...     # Solicita um ponto de carga
...     print('%s saindo em %d' % (name, env.now))
...     with bcs.request() as req:
...         yield req
...
...     # Carrega a bateria
...     print('%s carregando em %s' % (name, env.now))
...     yield env.timeout(charge_duration)
...     print('%s deixando a bcs em %s' % (name, env.now))
    
```

O método de Resource `request()` gera um evento que permite que você aguarde até que o recurso fique disponível novamente. Se você retornar, você “alocará” o recurso até que você o libere/*release*.

Caso utilize o recurso usando a declaração `with` como demonstrado a seguir, o recurso é automaticamente liberado. Porém se você chamar o `request()` sem declarar `with`, você toma pra si a responsabilidade de chamar o método `release()` quando você terminar de usar o recurso.

Quando liberar um recurso, o processo que estava aguardando por ele é continuado e passa a utilizar o recurso. A ordenação padrão de *Resource* na fila de espera segue o modelo de *FIFO* (*first in—first out*).

Um recurso precisa referenciar a classe *Environment* e ter uma capacidade/*capacity* determinada quando criado:

```

>>> import simpy
>>> env = simpy.Environment()
>>> bcs = simpy.Resource(env, capacity=2)
    
```

Tendo entendido essa primeira parte, podemos então criar os processos `car` e também informar uma referência para nosso *resource*/recurso como parâmetros adicionais:

```

>>> for i in range(4):
...     env.process(car(env, 'Car %d' % i, bcs, i*2, 5))
<Process(car) object at 0x...>
<Process(car) object at 0x...>
<Process(car) object at 0x...>
<Process(car) object at 0x...>
    
```

Então, finalmente podemos iniciar a simulação. Desde que os processos `car` Finally, we can start the simulation. Since the car processes all terminate on their own in this simulation, we don’t need to specify an *until* time—the simulation will automatically stop when there are no more events left:

```

>>> env.run()
Car 0 arriving at 0
Car 0 starting to charge at 0
Car 1 arriving at 2
Car 1 starting to charge at 2
Car 2 arriving at 4
Car 0 leaving the bcs at 5
Car 2 starting to charge at 5
Car 3 arriving at 6
Car 1 leaving the bcs at 7
Car 3 starting to charge at 7
Car 2 leaving the bcs at 10
Car 3 leaving the bcs at 12
    
```

Note that the first two cars can start charging immediately after they arrive at the BCS, while cars 2 an 3 have to wait.

## What's Next

You should now be familiar with SimPy's basic concepts. The [next section](#) shows you how you can proceed with using SimPy from here on.

## Como proceder

Caso ainda esteja em dúvida se o SimPy atende suas necessidade ou se pretende ver mais recursos em uso, você deveria analisar os diversos [exemplos](#) que nós disponibilizamos.

Se está a procura de uma descrição mais detalhada de um determinado aspecto ou recurso do Simpy, a seção [Topical Guides](#) pode lhe ajudar.

Finalmente, existe a [Documentação da API](#) que descreve todas as funções e classes em detalhes.

---

## Topical Guides

---

Esta seção abrange diversos aspectos do SimPy de forma mais aprofundada. Nela se assume que você tem um entendimento básico das capacidades do SimPy e que tem ideia do que está procurando.

### Conceitos básicos do SimPy

Este guia descreve os conceitos básicos do SimPy: Como funciona? O que são processos, eventos e o ambiente? O que pode ser feito com ele?

#### Como o SimPy funciona

Se você olhar melhor o SimPy, verá que ele é só um executor de eventos assíncronos. Você gera eventos e agenda eles para que sejam executados em um dado momento da simulação. Os Eventos são ordenados por prioridade, instante da simulação e um id incremental para cada evento registrado. Um evento também tem a uma lista de funções de retorno, os quais são executados quando o evento é acionado e processado em um loop. Os eventos podem retornar um valor.

The components involved in this are the *Environment*, *events* and the process functions that you write.

Process functions implement your simulation model, that is, they define the behavior of your simulation. They are plain Python generator functions that yield instances of *Event*.

The environment stores these events in its event list and keeps track of the current simulation time.

If a process function yields an event, SimPy adds the process to the event's callbacks and suspends the process until the event is triggered and processed. When a process waiting for an event is resumed, it will also receive the event's value.

Here is a very simple example that illustrates all this; the code is more verbose than it needs to be to make things extra clear. You find a compact version of it at the end of this section:

```
>>> import simpy
>>>
>>> def example(env):
...     event = simpy.events.Timeout(env, delay=1, value=42)
...     value = yield event
...     print('now=%d, value=%d' % (env.now, value))
>>>
>>> env = simpy.Environment()
>>> example_gen = example(env)
>>> p = simpy.events.Process(env, example_gen)
>>>
```

```
>>> env.run()
now=1, value=42
```

The `example()` process function above first creates a `Timeout` event. It passes the environment, a delay, and a value to it. The `Timeout` schedules itself at `now + delay` (that's why the environment is required); other event types usually schedule themselves at the current simulation time.

The process function then yields the event and thus gets suspended. It is resumed, when SimPy processes the `Timeout` event. The process function also receives the event's value (42) – this is, however, optional, so `yield event` would have been okay if the you were not interested in the value or if the event had no value at all.

Finally, the process function prints the current simulation time (that is accessible via the environment's `now` attribute) and the `Timeout`'s value.

If all required process functions are defined, you can instantiate all objects for your simulation. In most cases, you start by creating an instance of `Environment`, because you'll need to pass it around a lot when creating everything else.

Starting a process function involves two things:

1. You have to call the process function to create a generator object. (This will not execute any code of that function yet. Please read [The Python yield keyword explained](#), to understand why this is the case.)
2. You then create an instance of `Process` and pass the environment and the generator object to it. This will schedule an `Initialize` event at the current simulation time which starts the execution of the process function. The process instance is also an event that is triggered when the process function returns. The [guide to events](#) explains why this is handy.

Finally, you can start SimPy's event loop. By default, it will run as long as there are events in the event list, but you can also let it stop earlier by providing an `until` argument (see [Simulation control](#)).

The following guides describe the environment and its interactions with events and process functions in more detail.

## “Best practice” version of the example above

```
>>> import simpy
>>>
>>> def example(env):
...     value = yield env.timeout(1, value=42)
...     print('now=%d, value=%d' % (env.now, value))
>>>
>>> env = simpy.Environment()
>>> p = env.process(example(env))
>>> env.run()
now=1, value=42
```

## Environments

A simulation environment manages the simulation time as well as the scheduling and processing of events. It also provides means to step through or execute the simulation.

The base class for all environments is `BaseEnvironment`. “Normal” simulations usually use its subclass `Environment`. For real-time simulations, SimPy provides a `RealtimeEnvironment` (more on that in [Real-time simulations](#)).

## Simulation control

SimPy is very flexible in terms of simulation execution. You can run your simulation until there are no more events, until a certain simulation time is reached, or until a certain event is triggered. You can also step through the simulation event by event. Furthermore, you can mix these things as you like.

For example, you could run your simulation until an interesting event occurs. You could then step through the simulation event by event for a while; and finally run the simulation until there are no more events left and your processes have all terminated.

The most important method here is `Environment.run()`:

- If you call it without any argument (`env.run()`), it steps through the simulation until there are no more events left.

**Aviso:** If your processes run forever (`while True: yield env.timeout(1)`), this method will never terminate (unless you kill your script by e.g., pressing Ctrl-C).

- In most cases it is advisable to stop your simulation when it reaches a certain simulation time. Therefore, you can pass the desired time via the `until` parameter, e.g.: `env.run(until=10)`.

The simulation will then stop when the internal clock reaches 10 but will not process any events scheduled for time 10. This is similar to a new environment where the clock is 0 but (obviously) no events have yet been processed.

If you want to integrate your simulation in a GUI and want to draw a process bar, you can repeatedly call this function with increasing `until` values and update your progress bar after each call:

```
for i in range(100):
    env.run(until=i)
    progressbar.update(i)
```

- Instead of passing a number to `run()`, you can also pass any event to it. `run()` will then return when the event has been processed.

Assuming that the current time is 0, `env.run(until=env.timeout(5))` is equivalent to `env.run(until=5)`.

You can also pass other types of events (remember, that a `Process` is an event, too):

```
>>> import simpy
>>>
>>> def my_proc(env):
...     yield env.timeout(1)
...     return 'Monty Python's Flying Circus'
>>>
>>> env = simpy.Environment()
>>> proc = env.process(my_proc(env))
>>> env.run(until=proc)
'Monty Python's Flying Circus'
```

To step through the simulation event by event, the environment offers `peek()` and `step()`.

`peek()` returns the time of the next scheduled event or *infinity* (`float('inf')`) if no future events are scheduled.

`step()` processes the next scheduled event. It raises an `EmptySchedule` exception if no event is available.

In a typical use case, you use these methods in a loop like:

```
until = 10
while env.peek() < until:
    env.step()
```

## State access

The environment allows you to get the current simulation time via the `Environment.now` property. The simulation time is a number without unit and is increased via `Timeout` events.

By default, `now` starts at 0, but you can pass an `initial_time` to the `Environment` to use something else.

---

**Nota:** Although the simulation time is technically unitless, you can pretend that it is, for example, in seconds and use it like a timestamp returned by `time.time()` to calculate a date or the day of the week.

---

The property `Environment.active_process` is comparable to `os.getpid()` and is either `None` or pointing at the currently active `Process`. A process is *active* when its process function is being executed. It becomes *inactive* (or suspended) when it yields an event.

Thus, it only makes sense to access this property from within a process function or a function that is called by your process function:

```
>>> def subfunc(env):
...     print(env.active_process) # will print "p1"
>>>
>>> def my_proc(env):
...     while True:
...         print(env.active_process) # will print "p1"
...         subfunc(env)
...         yield env.timeout(1)
>>>
>>> env = simpy.Environment()
>>> p1 = env.process(my_proc(env))
>>> env.active_process # None
>>> env.step()
<Process(my_proc) object at 0x...>
<Process(my_proc) object at 0x...>
>>> env.active_process # None
```

An exemplary use case for this is the resource system: If a process function calls `request()` to request a resource, the resource determines the requesting process via `env.active_process`. Take a look at the code to see how we do this :-).

## Event creation

To create events, you normally have to import `simpy.events`, instantiate the event class and pass a reference to the environment to it. To reduce the amount of typing, the `Environment` provides some shortcuts for event creation. For example, `Environment.event()` is equivalent to `simpy.events.Event(env)`.

Other shortcuts are:

- `Environment.process()`
- `Environment.timeout()`
- `Environment.all_of()`

- `Environment.any_of()`

More details on what the events do can be found in the [guide to events](#).

## Miscellaneous

Since Python 3.3, a generator function can have a return value:

```
def my_proc(env):
    yield env.timeout(1)
    return 42
```

In SimPy, this can be used to provide return values for processes that can be used by other processes:

```
def other_proc(env):
    ret_val = yield env.process(my_proc(env))
    assert ret_val == 42
```

Internally, Python passes the return value as parameter to the `StopIteration` exception that it raises when a generator is exhausted. So in Python 2.7 and 3.2 you could replace the `return 42` with a `raise StopIteration(42)` to achieve the same result.

To keep your code more readable, the environment provides the method `exit()` to do exactly this:

```
def my_proc(env):
    yield env.timeout(1)
    env.exit(42) # Py2 equivalent to "return 42"
```

## Events

SimPy includes an extensive set of event types for various purposes. All of them inherit `simpy.events.Event`. The listing below shows the hierarchy of events built into SimPy:

```
events.Event
|
+-- events.Timeout
|
+-- events.Initialize
|
+-- events.Process
|
+-- events.Condition
| |
| +-- events.AllOf
| |
| +-- events.AnyOf
.
.
.
```

This is the set of basic events. Events are extensible and resources, for example, define additional events. In this guide, we'll focus on the events in the `simpy.events` module. The [guide to resources](#) describes the various resource events.

## Event basics

SimPy events are very similar – if not identical — to deferreds, futures or promises. Instances of the class `Event` are used to describe any kind of events. Events can be in one of the following states. An event

- might happen (not triggered),
- is going to happen (triggered) or
- has happened (processed).

They traverse these states exactly once in that order. Events are also tightly bound to time and time causes events to advance their state.

Initially, events are not triggered and just objects in memory.

If an event gets triggered, it is scheduled at a given time and inserted into SimPy's event queue. The property `Event.triggered` becomes `True`.

As long as the event is not *processed*, you can add *callbacks* to an event. Callbacks are callables that accept an event as parameter and are stored in the `Event.callbacks` list.

An event becomes *processed* when SimPy pops it from the event queue and calls all of its callbacks. It is now no longer possible to add callbacks. The property `Event.processed` becomes `True`.

Events also have a *value*. The value can be set before or when the event is triggered and can be retrieved via `Event.value` or, within a process, by yielding the event (`value = yield event`).

### Adding callbacks to an event

“What? Callbacks? I've never seen no callbacks!”, you might think if you have worked your way through the [tutorial](#).

That's on purpose. The most common way to add a callback to an event is yielding it from your process function (`yield event`). This will add the process' `_resume()` method as a callback. That's how your process gets resumed when it yielded an event.

However, you can add any callable object (function) to the list of callbacks as long as it accepts an event instance as its single parameter:

```
>>> import simpy
>>>
>>> def my_callback(event):
...     print('Called back from', event)
...
>>> env = simpy.Environment()
>>> event = env.event()
>>> event.callbacks.append(my_callback)
>>> event.callbacks
[<function my_callback at 0x...>]
```

If an event has been *processed*, all of its `Event.callbacks` have been executed and the attribute is set to `None`. This is to prevent you from adding more callbacks – these would of course never get called because the event has already happened.

Processes are smart about this, though. If you yield a processed event, `_resume()` will immediately resume your process with the value of the event (because there is nothing to wait for).

## Triggering events

When events are triggered, they can either *succeed* or *fail*. For example, if an event is to be triggered at the end of a computation and everything works out fine, the event will *succeed*. If an exceptions occurs during that computation, the event will *fail*.

To trigger an event and mark it as successful, you can use `Event.succeed(value=None)`. You can optionally pass a *value* to it (e.g., the results of a computation).

To trigger an event and mark it as failed, call `Event.fail(exception)` and pass an `Exception` instance to it (e.g., the exception you caught during your failed computation).

There is also a generic way to trigger an event: `Event.trigger(event)`. This will take the value and outcome (success or failure) of the event passed to it.

All three methods return the event instance they are bound to. This allows you to do things like `yield Event(env).succeed()`.

## Example usages for Event

The simple mechanics outlined above provide a great flexibility in the way events (even the basic `Event`) can be used.

One example for this is that events can be shared. They can be created by a process or outside of the context of a process. They can be passed to other processes and chained:

```
>>> class School:
...     def __init__(self, env):
...         self.env = env
...         self.class_ends = env.event()
...         self.pupil_procs = [env.process(self.pupil()) for i in range(3)]
...         self.bell_proc = env.process(self.bell())
...
...     def bell(self):
...         for i in range(2):
...             yield self.env.timeout(45)
...             self.class_ends.succeed()
...             self.class_ends = self.env.event()
...             print()
...
...     def pupil(self):
...         for i in range(2):
...             print(' \o/', end='')
...             yield self.class_ends
...
>>> school = School(env)
>>> env.run()
\o/ \o/ \o/
\o/ \o/ \o/
```

This can also be used like the *passivate / reactivate* known from SimPy 2. The pupils *passivate* when class begins and are *reactivated* when the bell rings.

## Let time pass by: the Timeout

To actually let time pass in a simulation, there is the *timeout* event. A timeout has two parameters: a *delay* and an optional *value*: `Timeout(delay, value=None)`. It triggers itself during its creation and schedules itself at now

+ delay. Thus, the `succeed()` and `fail()` methods cannot be called again and you have to pass the event value to it when you create the timeout.

The delay can be any kind of number, usually an *int* or *float* as long as it supports comparison and addition.

## Processes are events, too

SimPy processes (as created by `Process` or `env.process()`) have the nice property of being events, too.

That means, that a process can yield another process. It will then be resumed when the other process ends. The event's value will be the return value of that process:

```
>>> def sub(env):
...     yield env.timeout(1)
...     return 23
...
>>> def parent(env):
...     ret = yield env.process(sub(env))
...     return ret
...
>>> env.run(env.process(parent(env)))
23
```

The example above will only work in Python  $\geq 3.3$ . As a workaround for older Python versions, you can use `env.exit(23)` with the same effect.

When a process is created, it schedules an *Initialize* event which will start the execution of the process when triggered. You usually won't have to deal with this type of event.

If you don't want a process to start immediately but after a certain delay, you can use `simpy.util.start_delayed()`. This method returns a helper process that uses a *timeout* before actually starting a process.

The example from above, but with a delayed start of `sub()`:

```
>>> from simpy.util import start_delayed
>>>
>>> def sub(env):
...     yield env.timeout(1)
...     return 23
...
>>> def parent(env):
...     start = env.now
...     sub_proc = yield start_delayed(env, sub(env), delay=3)
...     assert env.now - start == 3
...
...     ret = yield sub_proc
...     return ret
...
>>> env.run(env.process(parent(env)))
23
```

## Waiting for multiple events at once

Sometimes, you want to wait for more than one event at the same time. For example, you may want to wait for a resource, but not for an unlimited amount of time. Or you may want to wait until all a set of events has happened.

SimPy therefore offers the *AnyOf* and *AllOf* events which both are a *Condition* event.

Both take a list of events as an argument and are triggered if at least one or all of them are triggered.

```
>>> from simpy.events import AnyOf, AllOf, Event
>>> events = [Event(env) for i in range(3)]
>>> a = AnyOf(env, events) # Triggers if at least one of "events" is triggered.
>>> b = AllOf(env, events) # Triggers if all each of "events" is triggered.
```

The value of a condition event is an ordered dictionary with an entry for every triggered event. In the case of `AllOf`, the size of that dictionary will always be the same as the length of the event list. The value dict of `AnyOf` will have at least one entry. In both cases, the event instances are used as keys and the event values will be the values.

As a shorthand for `AllOf` and `AnyOf`, you can also use the logical operators `&` (and) and `|` (or):

```
>>> def test_condition(env):
...     t1, t2 = env.timeout(1, value='spam'), env.timeout(2, value='eggs')
...     ret = yield t1 | t2
...     assert ret == {t1: 'spam'}
...
...     t1, t2 = env.timeout(1, value='spam'), env.timeout(2, value='eggs')
...     ret = yield t1 & t2
...     assert ret == {t1: 'spam', t2: 'eggs'}
...
...     # You can also concatenate & and |
...     e1, e2, e3 = [env.timeout(i) for i in range(3)]
...     yield (e1 | e2) & e3
...     assert all(e.processed for e in [e1, e2, e3])
...
>>> proc = env.process(test_condition(env))
>>> env.run()
```

The order of condition results is identical to the order in which the condition events were specified. This allows the following idiom for conveniently fetching the values of multiple events specified in an *and* condition (including `AllOf`):

```
>>> def fetch_values_of_multiple_events(env):
...     t1, t2 = env.timeout(1, value='spam'), env.timeout(2, value='eggs')
...     r1, r2 = (yield t1 & t2).values()
...     assert r1 == 'spam' and r2 == 'eggs'
...
>>> proc = env.process(fetch_values_of_multiple_events(env))
>>> env.run()
```

## Process Interaction

Discrete event simulation is only made interesting by interactions between processes.

So this guide is about:

- *Sleep until woken up* (passivate/reactivate)
- *Waiting for another process to terminate*
- *Interrupting another process*

The first two items were already covered in the [Events](#) guide, but we'll also include them here for the sake of completeness.

Another possibility for processes to interact are resources. They are discussed in a [separate guide](#).

## Sleep until woken up

Imagine you want to model an electric vehicle with an intelligent battery-charging controller. While the vehicle is driving, the controller can be passive but needs to be reactivated once the vehicle is connected to the power grid in order to charge the battery.

In SimPy 2, this pattern was known as *passivate / reactivate*. In SimPy 3, you can accomplish that with a simple, shared *Event*:

```
>>> from random import seed, randint
>>> seed(23)
>>>
>>> import simpy
>>>
>>> class EV:
...     def __init__(self, env):
...         self.env = env
...         self.drive_proc = env.process(self.drive(env))
...         self.bat_ctrl_proc = env.process(self.bat_ctrl(env))
...         self.bat_ctrl_reactivate = env.event()
...
...     def drive(self, env):
...         while True:
...             # Drive for 20-40 min
...             yield env.timeout(randint(20, 40))
...
...             # Park for 1-6 hours
...             print('Start parking at', env.now)
...             self.bat_ctrl_reactivate.succeed() # "reactivate"
...             self.bat_ctrl_reactivate = env.event()
...             yield env.timeout(randint(60, 360))
...             print('Stop parking at', env.now)
...
...     def bat_ctrl(self, env):
...         while True:
...             print('Bat. ctrl. passivating at', env.now)
...             yield self.bat_ctrl_reactivate # "passivate"
...             print('Bat. ctrl. reactivated at', env.now)
...
...             # Intelligent charging behavior here ...
...             yield env.timeout(randint(30, 90))
...
>>> env = simpy.Environment()
>>> ev = EV(env)
>>> env.run(until=150)
Bat. ctrl. passivating at 0
Start parking at 29
Bat. ctrl. reactivated at 29
Bat. ctrl. passivating at 60
Stop parking at 131
```

Since `bat_ctrl()` just waits for a normal event, we no longer call this pattern *passivate / reactivate* in SimPy 3.

## Waiting for another process to terminate

The example above has a problem: it may happen that the vehicles wants to park for a shorter duration than it takes to charge the battery (this is the case if both, charging and parking would take 60 to 90 minutes).

To fix this problem we have to slightly change our model. A new `bat_ctrl()` will be started every time the EV starts parking. The EV then waits until the parking duration is over *and* until the charging has stopped:

```
>>> class EV:
...     def __init__(self, env):
...         self.env = env
...         self.drive_proc = env.process(self.drive(env))
...
...     def drive(self, env):
...         while True:
...             # Drive for 20-40 min
...             yield env.timeout(randint(20, 40))
...
...             # Park for 1-6 hours
...             print('Start parking at', env.now)
...             charging = env.process(self.bat_ctrl(env))
...             parking = env.timeout(randint(60, 360))
...             yield charging & parking
...             print('Stop parking at', env.now)
...
...     def bat_ctrl(self, env):
...         print('Bat. ctrl. started at', env.now)
...         # Intelligent charging behavior here ...
...         yield env.timeout(randint(30, 90))
...         print('Bat. ctrl. done at', env.now)
...
>>> env = simpy.Environment()
>>> ev = EV(env)
>>> env.run(until=310)
Start parking at 29
Bat. ctrl. started at 29
Bat. ctrl. done at 83
Stop parking at 305
```

Again, nothing new (if you've read the [Events](#) guide) and special is happening. SimPy processes are events, too, so you can yield them and will thus wait for them to get triggered. You can also wait for two events at the same time by concatenating them with `&` (see [Waiting for multiple events at once](#)).

## Interrupting another process

As usual, we now have another problem: Imagine, a trip is very urgent, but with the current implementation, we always need to wait until the battery is fully charged. If we could somehow interrupt that ...

Fortunate coincidence, there is indeed a way to do exactly this. You can call `interrupt()` on a *Process*. This will throw an *Interrupt* exception into that process, resuming it immediately:

```
>>> class EV:
...     def __init__(self, env):
...         self.env = env
...         self.drive_proc = env.process(self.drive(env))
...
...     def drive(self, env):
...         while True:
...             # Drive for 20-40 min
...             yield env.timeout(randint(20, 40))
...
...             # Park for 1 hour
...             print('Start parking at', env.now)
```

```

...         charging = env.process(self.bat_ctrl(env))
...         parking = env.timeout(60)
...         yield charging | parking
...         if not charging.triggered:
...             # Interrupt charging if not already done.
...             charging.interrupt('Need to go!')
...         print('Stop parking at', env.now)
...
...     def bat_ctrl(self, env):
...         print('Bat. ctrl. started at', env.now)
...         try:
...             yield env.timeout(randint(60, 90))
...             print('Bat. ctrl. done at', env.now)
...         except simpy.Interrupt as i:
...             # Onoes! Got interrupted before the charging was done.
...             print('Bat. ctrl. interrupted at', env.now, 'msg:',
...                   i.cause)
...
>>> env = simpy.Environment()
>>> ev = EV(env)
>>> env.run(until=100)
Start parking at 31
Bat. ctrl. started at 31
Stop parking at 91
Bat. ctrl. interrupted at 91 msg: Need to go!

```

What `process.interrupt()` actually does is scheduling an *Interruption* event for immediate execution. If this event is executed it will remove the victim process' `_resume()` method from the callbacks of the event that it is currently waiting for (see *target*). Following that it will throw the `Interrupt` exception into the process.

Since we don't do anything special to the original target event of the process, the interrupted process can yield the same event again after catching the `Interrupt` – Imagine someone waiting for a shop to open. The person may get interrupted by a phone call. After finishing the call, he or she checks if the shop already opened and either enters or continues to wait.

## Shared Resources

Shared resources are another way to model *Process Interaction*. They form a congestion point where processes queue up in order to use them.

SimPy defines three categories of resources:

- *Resources* – Resources that can be used by a limited number of processes at a time (e.g., a gas station with a limited number of fuel pumps).
- *Containers* – Resources that model the production and consumption of a homogeneous, undifferentiated bulk. It may either be continuous (like water) or discrete (like apples).
- *Stores* – Resources that allow the production and consumption of Python objects.

## The basic concept of resources

All resources share the same basic concept: The resource itself is some kind of a container with a, usually limited, *capacity*. Processes can either try to *put* something into the resource or try to *get* something out. If the resource is full or empty, they have to *queue* up and wait.

This is roughly how every resource looks:

```
BaseResource(capacity):
    put_queue
    get_queue

    put(): event
    get(): event
```

Every resource has a maximum capacity and two queues: one for processes that want to put something into it and one for processes that want to get something out. The `put()` and `get()` methods both return an event that is triggered when the corresponding action was successful.

## Resources and interrupts

While a process is waiting for a put or get event to succeed, it may be *interrupted* by another process. After catching the interrupt, the process has two possibilities:

1. It may continue to wait for the request (by yielding the event again).
2. It may stop waiting for the request. In this case, it has to call the event's `cancel()` method.

Since you can easily forget this, all resources events are *context managers* (see the [Python docs](#) for details).

The resource system is modular and extensible. Resources can, for example, use specialized queues and event types. This allows them to use sorted queues, to add priorities to events, or to offer preemption.

## Resources

Resources can be used by a limited number of processes at a time (e.g., a gas station with a limited number of fuel pumps). Processes *request* these resources to become a user (or to “own” them) and have to *release* them once they are done (e.g., vehicles arrive at the gas station, use a fuel-pump, if one is available, and leave when they are done).

Requesting a resource is modeled as “putting a process’ token into the resource” and releasing a resource correspondingly as “getting a process’ token out of the resource”. Thus, calling `request()/release()` is equivalent to calling `put()/get()`. Releasing a resource will always succeed immediately.

SimPy implements three *resource* types:

1. *Resource*
2. *PriorityResource*, where queueing processes are sorted by priority
3. *PreemptiveResource*, where processes additionally may preempt other processes with a lower priority

## Resource

The *Resource* is conceptually a *semaphore*. Its only parameter – apart from the obligatory reference to an *Environment* – is its *capacity*. It must be a positive number and defaults to 1: `Resource(env, capacity=1)`.

Instead of just counting its current users, it stores the request event as an “access token” for each user. This is, for example, useful for adding preemption (see below).

Here is a basic example for using a resource:

```

>>> import simpy
>>>
>>> def resource_user(env, resource):
...     request = resource.request() # Generate a request event
...     yield request                # Wait for access
...     yield env.timeout(1)         # Do something
...     resource.release(request)    # Release the resource
...
>>> env = simpy.Environment()
>>> res = simpy.Resource(env, capacity=1)
>>> user = env.process(resource_user(env, res))
>>> env.run()
    
```

Note, that you have to release the resource under all conditions; for example, if you got interrupted while waiting for or using the resource. In order to help you with that and to avoid too many `try: ... finally: ...` constructs, request events can be used as context manager:

```

>>> def resource_user(env, resource):
...     with resource.request() as req: # Generate a request event
...         yield req                  # Wait for access
...         yield env.timeout(1)       # Do something
...                                     # Resource released automatically
>>> user = env.process(resource_user(env, res))
>>> env.run()
    
```

Resources allow you to retrieve lists of the current users or queued users, the number of current users and the resource's capacity:

```

>>> res = simpy.Resource(env, capacity=1)
>>>
>>> def print_stats(res):
...     print('%d of %d slots are allocated.' % (res.count, res.capacity))
...     print('  Users:', res.users)
...     print('  Queued events:', res.queue)
>>>
>>>
>>> def user(res):
...     print_stats(res)
...     with res.request() as req:
...         yield req
...         print_stats(res)
...     print_stats(res)
>>>
>>> procs = [env.process(user(res)), env.process(user(res))]
>>> env.run()
0 of 1 slots are allocated.
  Users: []
  Queued events: []
1 of 1 slots are allocated.
  Users: [<Request() object at 0x...>]
  Queued events: []
1 of 1 slots are allocated.
  Users: [<Request() object at 0x...>]
  Queued events: [<Request() object at 0x...>]
0 of 1 slots are allocated.
  Users: []
  Queued events: [<Request() object at 0x...>]
1 of 1 slots are allocated.
  Users: [<Request() object at 0x...>]
    
```

```

Queued events: []
0 of 1 slots are allocated.
Users: []
Queued events: []

```

## PriorityResource

As you may know from the real world, not every one is equally important. To map that to SimPy, there's the *PriorityResource*. This subclass of *Resource* lets requesting processes provide a priority for each request. More important requests will gain access to the resource earlier than less important ones. Priority is expressed by integer numbers; smaller numbers mean a higher priority.

Apart from that, it works like a normal *Resource*:

```

>>> def resource_user(name, env, resource, wait, prio):
...     yield env.timeout(wait)
...     with resource.request(priority=prio) as req:
...         print('%s requesting at %s with priority=%s' % (name, env.now, prio))
...         yield req
...         print('%s got resource at %s' % (name, env.now))
...         yield env.timeout(3)
...
>>> env = simpy.Environment()
>>> res = simpy.PriorityResource(env, capacity=1)
>>> p1 = env.process(resource_user(1, env, res, wait=0, prio=0))
>>> p2 = env.process(resource_user(2, env, res, wait=1, prio=0))
>>> p3 = env.process(resource_user(3, env, res, wait=2, prio=-1))
>>> env.run()
1 requesting at 0 with priority=0
1 got resource at 0
2 requesting at 1 with priority=0
3 requesting at 2 with priority=-1
3 got resource at 3
2 got resource at 6

```

Although *p3* requested the resource later than *p2*, it could use it earlier because its priority was higher.

## PreemptiveResource

Sometimes, new requests are so important that queue-jumping is not enough and they need to kick existing users out of the resource (this is called *preemption*). The *PreemptiveResource* allows you to do exactly this:

```

>>> def resource_user(name, env, resource, wait, prio):
...     yield env.timeout(wait)
...     with resource.request(priority=prio) as req:
...         print('%s requesting at %s with priority=%s' % (name, env.now, prio))
...         yield req
...         print('%s got resource at %s' % (name, env.now))
...         try:
...             yield env.timeout(3)
...         except simpy.Interrupt as interrupt:
...             by = interrupt.cause.by
...             usage = env.now - interrupt.cause.usage_since
...             print('%s got preempted by %s at %s after %s' %
...                   (name, by, env.now, usage))
...

```

```

>>> env = simpy.Environment()
>>> res = simpy.PreemptiveResource(env, capacity=1)
>>> p1 = env.process(resource_user(1, env, res, wait=0, prio=0))
>>> p2 = env.process(resource_user(2, env, res, wait=1, prio=0))
>>> p3 = env.process(resource_user(3, env, res, wait=2, prio=-1))
>>> env.run()
1 requesting at 0 with priority=0
1 got resource at 0
2 requesting at 1 with priority=0
3 requesting at 2 with priority=-1
1 got preempted by <Process(resource_user) object at 0x...> at 2 after 2
3 got resource at 2
2 got resource at 5
    
```

*PreemptiveResource* inherits from *PriorityResource* and adds a *preempt* flag (that defaults to `True`) to `request()`. By setting this to `False` (`resource.request(priority=x, preempt=False)`), a process can decide to not preempt another resource user. It will still be put in the queue according to its priority, though.

The implementation of *PreemptiveResource* values priorities higher than *preemption*. That means preempt requests are not allowed to cheat and jump over a higher prioritized request. The following example shows that preemptive low priority requests cannot queue-jump over high priority requests:

```

>>> def user(name, env, res, prio, preempt):
...     with res.request(priority=prio, preempt=preempt) as req:
...         try:
...             print('%s requesting at %d' % (name, env.now))
...             yield req
...             print('%s got resource at %d' % (name, env.now))
...             yield env.timeout(3)
...         except simpy.Interrupt:
...             print('%s got preempted at %d' % (name, env.now))
>>>
>>> env = simpy.Environment()
>>> res = simpy.PreemptiveResource(env, capacity=1)
>>> A = env.process(user('A', env, res, prio=0, preempt=True))
>>> env.run(until=1) # Give A a head start
A requesting at 0
A got resource at 0
>>> B = env.process(user('B', env, res, prio=-2, preempt=False))
>>> C = env.process(user('C', env, res, prio=-1, preempt=True))
>>> env.run()
B requesting at 1
C requesting at 1
B got resource at 3
C got resource at 6
    
```

1. Process *A* requests the resource with priority 0. It immediately becomes a user.
2. Process *B* requests the resource with priority -2 but sets *preempt* to `False`. It will queue up and wait.
3. Process *C* requests the resource with priority -1 but leaves *preempt* `True`. Normally, it would preempt *A* but in this case, *B* is queued up before *C* and prevents *C* from preempting *A*. *C* can also not preempt *B* since its priority is not high enough.

Thus, the behavior in the example is the same as if no preemption was used at all. Be careful when using mixed preemption!

Due to the higher priority of process *B*, no preemption occurs in this example. Note that an additional request with a priority of -3 would be able to preempt *A*.

If your use-case requires a different behaviour, for example queue-jumping or valuing preemption over priorities, you can subclass *PreemptiveResource* and override the default behaviour.

## Containers

Containers help you modelling the production and consumption of a homogeneous, undifferentiated bulk. It may either be continuous (like water) or discrete (like apples).

You can use this, for example, to model the gas / petrol tank of a gas station. Tankers increase the amount of gasoline in the tank while cars decrease it.

The following example is a very simple model of a gas station with a limited number of fuel dispensers (modeled as Resource) and a tank modeled as Container:

```
>>> class GasStation:
...     def __init__(self, env):
...         self.fuel_dispensers = simpy.Resource(env, capacity=2)
...         self.gas_tank = simpy.Container(env, init=100, capacity=1000)
...         self.mon_proc = env.process(self.monitor_tank(env))
...
...     def monitor_tank(self, env):
...         while True:
...             if self.gas_tank.level < 100:
...                 print('Calling tanker at %s' % env.now)
...                 env.process(tanker(env, self))
...             yield env.timeout(15)
>>>
>>>
>>> def tanker(env, gas_station):
...     yield env.timeout(10) # Need 10 Minutes to arrive
...     print('Tanker arriving at %s' % env.now)
...     amount = gas_station.gas_tank.capacity - gas_station.gas_tank.level
...     yield gas_station.gas_tank.put(amount)
>>>
>>>
>>> def car(name, env, gas_station):
...     print('Car %s arriving at %s' % (name, env.now))
...     with gas_station.fuel_dispensers.request() as req:
...         yield req
...         print('Car %s starts refueling at %s' % (name, env.now))
...         yield gas_station.gas_tank.get(40)
...         yield env.timeout(5)
...         print('Car %s done refueling at %s' % (name, env.now))
>>>
>>>
>>> def car_generator(env, gas_station):
...     for i in range(4):
...         env.process(car(i, env, gas_station))
...         yield env.timeout(5)
>>>
>>>
>>> env = simpy.Environment()
>>> gas_station = GasStation(env)
>>> car_gen = env.process(car_generator(env, gas_station))
>>> env.run(35)
Car 0 arriving at 0
Car 0 starts refueling at 0
Car 1 arriving at 5
```

```

Car 0 done refueling at 5
Car 1 starts refueling at 5
Car 2 arriving at 10
Car 1 done refueling at 10
Car 2 starts refueling at 10
Calling tanker at 15
Car 3 arriving at 15
Car 3 starts refueling at 15
Tanker arriving at 25
Car 2 done refueling at 30
Car 3 done refueling at 30

```

Containers allow you to retrieve their current level as well as their capacity (see `GasStation.monitor_tank()` and `tanker()`). You can also access the list of waiting events via the `put_queue` and `get_queue` attributes (similar to `Resource.queue`).

## Stores

Using Stores you can model the production and consumption of concrete objects (in contrast to the rather abstract “amount” stored in containers). A single Store can even contain multiple types of objects.

Beside *Store*, there is a *FilterStore* that lets you use a custom function to filter the objects you get out of the store and *PriorityStore* where items come out of the store in priority order.

Here is a simple example modelling a generic producer/consumer scenario:

```

>>> def producer(env, store):
...     for i in range(100):
...         yield env.timeout(2)
...         yield store.put('spam %s' % i)
...         print('Produced spam at', env.now)
>>>
>>>
>>> def consumer(name, env, store):
...     while True:
...         yield env.timeout(1)
...         print(name, 'requesting spam at', env.now)
...         item = yield store.get()
...         print(name, 'got', item, 'at', env.now)
>>>
>>>
>>> env = simpy.Environment()
>>> store = simpy.Store(env, capacity=2)
>>>
>>> prod = env.process(producer(env, store))
>>> consumers = [env.process(consumer(i, env, store)) for i in range(2)]
>>>
>>> env.run(until=5)
0 requesting spam at 1
1 requesting spam at 1
Produced spam at 2
0 got spam 0 at 2
0 requesting spam at 3
Produced spam at 4
1 got spam 1 at 4

```

As with the other resource types, you can get a store’s capacity via the `capacity` attribute. The attribute `items` points to the list of items currently available in the store. The `put` and `get` queues can be accessed via the `put_queue`

and `get_queue` attributes.

`FilterStore` can, for example, be used to model machine shops where machines have varying attributes. This can be useful if the homogeneous slots of a `Resource` are not what you need:

```
>>> from collections import namedtuple
>>>
>>> Machine = namedtuple('Machine', 'size, duration')
>>> m1 = Machine(1, 2) # Small and slow
>>> m2 = Machine(2, 1) # Big and fast
>>>
>>> env = simpy.Environment()
>>> machine_shop = simpy.FilterStore(env, capacity=2)
>>> machine_shop.items = [m1, m2] # Pre-populate the machine shop
>>>
>>> def user(name, env, ms, size):
...     machine = yield ms.get(lambda machine: machine.size == size)
...     print(name, 'got', machine, 'at', env.now)
...     yield env.timeout(machine.duration)
...     yield ms.put(machine)
...     print(name, 'released', machine, 'at', env.now)
>>>
>>>
>>> users = [env.process(user(i, env, machine_shop, (i % 2) + 1))
...          for i in range(3)]
>>> env.run()
0 got Machine(size=1, duration=2) at 0
1 got Machine(size=2, duration=1) at 0
1 released Machine(size=2, duration=1) at 1
0 released Machine(size=1, duration=2) at 2
2 got Machine(size=1, duration=2) at 2
2 released Machine(size=1, duration=2) at 4
```

With a `PriorityStore`, we can model items of differing priorities. In the following example, an inspector process finds and logs issues that a separate maintainer process repairs in priority order.

```
>>> env = simpy.Environment()
>>> issues = simpy.PriorityStore(env)
>>>
>>> def inspector(env, issues):
...     for issue in [simpy.PriorityItem('P2', '#0000'),
...                  simpy.PriorityItem('P0', '#0001'),
...                  simpy.PriorityItem('P3', '#0002'),
...                  simpy.PriorityItem('P1', '#0003')]:
...         yield env.timeout(1)
...         print(env.now, 'log', issue)
...         yield issues.put(issue)
>>>
>>> def maintainer(env, issues):
...     while True:
...         yield env.timeout(3)
...         issue = yield issues.get()
...         print(env.now, 'repair', issue)
>>>
>>> _ = env.process(inspector(env, issues))
>>> _ = env.process(maintainer(env, issues))
>>> env.run()
1 log PriorityItem(priority='P2', item='#0000')
2 log PriorityItem(priority='P0', item='#0001')
```

```

3 log PriorityItem(priority='P3', item='#0002')
3 repair PriorityItem(priority='P0', item='#0001')
4 log PriorityItem(priority='P1', item='#0003')
6 repair PriorityItem(priority='P1', item='#0003')
9 repair PriorityItem(priority='P2', item='#0000')
12 repair PriorityItem(priority='P3', item='#0002')

```

## Real-time simulations

Sometimes, you might not want to perform a simulation as fast as possible but synchronous to the wall-clock time. This kind of simulation is also called *real-time simulation*.

Real-time simulations may be necessary

- if you have hardware-in-the-loop,
- if there is human interaction with your simulation, or
- if you want to analyze the real-time behavior of an algorithm.

To convert a simulation into a real-time simulation, you only need to replace SimPy's default *Environment* with a *simpy.rt.RealtimeEnvironment*. Apart from the *initial\_time* argument, there are two additional parameters: *factor* and *strict*: `RealtimeEnvironment(initial_time=0, factor=1.0, strict=True)`.

The *factor* defines how much *real time* passes with each step of simulation time. By default, this is one second. If you set `factor=0.1`, a unit of simulation time will only take a tenth of a second; if you set `factor=60`, it will take a minute.

Here is a simple example for converting a normal simulation to a real-time simulation with a duration of one tenth of a second per simulation time unit:

```

>>> import time
>>> import simpy
>>>
>>> def example(env):
...     start = time.perf_counter()
...     yield env.timeout(1)
...     end = time.perf_counter()
...     print('Duration of one simulation time unit: %.2fs' % (end - start))
>>>
>>> env = simpy.Environment()
>>> proc = env.process(example(env))
>>> env.run(until=proc)
Duration of one simulation time unit: 0.00s
>>>
>>> import simpy.rt
>>> env = simpy.rt.RealtimeEnvironment(factor=0.1)
>>> proc = env.process(example(env))
>>> env.run(until=proc)
Duration of one simulation time unit: 0.10s

```

If the *strict* parameter is set to `True` (the default), the `step()` and `run()` methods will raise a `RuntimeError` if the computation within a simulation time step take more time than the real-time factor allows. In the following example, a process will perform a task that takes 0.02 seconds within a real-time environment with a time factor of 0.01 seconds:

```

>>> import time
>>> import simpy.rt

```

```

>>>
>>> def slow_proc(env):
...     time.sleep(0.02) # Heavy computation :-
...     yield env.timeout(1)
>>>
>>> env = simpy.rt.RealtimeEnvironment(factor=0.01)
>>> proc = env.process(slow_proc(env))
>>> try:
...     env.run(until=proc)
...     print('Everything alright')
... except RuntimeError:
...     print('Simulation is too slow')
Simulation is too slow

```

To suppress the error, simply set `strict=False`:

```

>>> env = simpy.rt.RealtimeEnvironment(factor=0.01, strict=False)
>>> proc = env.process(slow_proc(env))
>>> try:
...     env.run(until=proc)
...     print('Everything alright')
... except RuntimeError:
...     print('Simulation is too slow')
Everything alright

```

That's it. Real-time simulations are that simple with SimPy!

## Monitorando

O processo de monitoramento é um tópico relativamente complexo por ter diferentes casos de uso e uma boa diversidade de variações.

Este guia apresenta alguns dos casos mais comuns e possivelmente que mais se destacam. São propostas que podem lhe dar algumas dicas/idéias de como implementar o monitoramento de uma simulação e adaptá-la a sua necessidade.

Porém, antes de começar é preciso definir:

*O que* você deseja monitorar?

- *Seus Processos?*
- *Uso de recursos?*
- *Rastrear todos os Eventos da simulação?*

*Quando* você quer monitorar?

- Em intervalos regulares?
- Quando algo ocorrer?

*Como* pretende armazenar os dados coletados?

- Em uma lista simples?
- Registrar em um arquivo?
- Armazenar em uma base de dados?

As seções a seguir analisam estas perguntas e provêem alguns exemplos para lhe ajudar.

## Monitorando seus processos

Monitorar seus próprios processos é uma coisa relativamente simples, afinal *you* controla o código. Temos conhecimento que a coisa mais comum que se deseja fazer é monitorar o valor de uma ou mais variáveis de estado toda vez que elas mudarem ou em um intervalo de tempo e armazenar essa informação em algum lugar ( na memória, em um banco de dados ou até mesmo um arquivo, por exemplo).

Um exemplo simples seria o uso de uma lista e adicionar os valores desejados para análise a cada vez que ela mudasse:

```
>>> import simpy
>>>
>>> data = [] # Esta lista armazenara os dados coletados
>>>
>>> def test_process(env, data):
...     val = 0
...     for i in range(5):
...         val += env.now
...         data.append(val) # Armazenando o dado
...         yield env.timeout(1)
>>>
>>> env = simpy.Environment()
>>> p = env.process(test_process(env, data))
>>> env.run(p)
>>> print('Collected', data) # Imprimindo os dados coletados
Collected [0, 1, 3, 6, 10]
```

Se precisa monitorar múltiplas variáveis, você pode inserir tuplas em sua lista.

Caso queira armazenar os dados em um array Numpy ou uma base de dados, muitas vezes você pode melhorar o desempenho se você guardar os dados em uma lista no Python e só escrever as partes maiores ( ou o dataset completo ) em uma base de dados.

## Uso de recursos

São muitos os casos de uso para monitoramento de recursos, por exemplo você pode desejar monitorar:

- O uso de um recurso no decorrer da simulação e ter média de uso, que seria,
  - o número de processos que usaram o recurso no decorrer da simulação
  - o nível de um contêiner
  - a quantidade de itens em uma loja

Isto pode ser monitorado também no decorrer do tempo da simulação discreta ou a cada vez que houver uma mudança.

- O número de processos armazenados ou retirados de uma fila no decorrer da simulação ( e a quantidade média ). Mais uma vez, isto pode ser monitorado no decorrer do tempo da simulação ou a cada mudança.
- Para *Recursos Preemptivos*, você pode desejar medir a frequência com que a preempção ocorre no decorrer do tempo da simulação.

Ao contrário do que ocorre nos processos, você não tem acesso direto ao código das classes de recurso embutidas. Porém isso não o impede de monitorá-las.

Algumas adaptações nos métodos dos recursos permitem que você capture todos os dados necessários.

Segue um exemplo que demonstra como você pode adicionar retornos para um recurso que deverá ser chamado apenas antes ou depois de um evento *get / request* ou um *put / release*:

```

>>> from functools import partial, wraps
>>> import simpy
>>>
>>> def patch_resource(resource, pre=None, post=None):
...     """Ajuste Patch *resource* que executa o metodo *pre* antes de cada
...     operacao de put/get/request/release e o metodo *post* executado apos
...     cada operacao. O unico argumento solicitado para essas funcoes eh a
...     instancia do recurso.
...
...     """
...     def get_wrapper(func):
...         # Cria um wrapper para put/get/request/release
...         @wraps(func)
...         def wrapper(*args, **kwargs):
...             # Este wrapper executa a
...             # funcao "pre"
...             if pre:
...                 pre(resource)
...
...             # Executa a operacao atual
...             ret = func(*args, **kwargs)
...
...             # Chama a funcao "post"
...             if post:
...                 post(resource)
...
...             return ret
...         return wrapper
...
...     # Substitui as operacoes originais com nosso wrapper
...     for name in ['put', 'get', 'request', 'release']:
...         if hasattr(resource, name):
...             setattr(resource, name, get_wrapper(getattr(resource, name)))
>>>
>>> def monitor(data, resource):
...     """Esta eh a funcao de monitoramento."""
...     item = (
...         resource._env.now, # O tempo na simulacao atual
...         resource.count, # O numero de usuarios
...         len(resource.queue), # O numero de processos enfileirados
...     )
...     data.append(item)
>>>
>>> def test_process(env, res):
...     with res.request() as req:
...         yield req
...         yield env.timeout(1)
>>>
>>> env = simpy.Environment()
>>>
>>> res = simpy.Resource(env, capacity=1)
>>> data = []
>>> # Vincular *data* como o primeiro argumento de monitor()
>>> # veja mais em https://docs.python.org/3/library/functools.html#functools.partial
>>> monitor = partial(monitor, data)
>>> patch_resource(res, post=monitor) # Altera/Modifica (somente) esta instancia do recurso
>>>
>>> p = env.process(test_process(env, res))

```

```
>>> env.run(p)
>>>
>>> print(data)
[(0, 1, 0), (1, 0, 0)]
```

O exemplo acima é bem genérico mas também é uma forma bem flexível de monitorar todos os aspectos de diversos tipos de recursos.

O outro extremo seria tornar o monitoramento um caso de uso. Suponha que você só quer saber quantos processos estão aguardando por um `Resource` em um determinado momento:

```
>>> import simpy
>>>
>>> class MonitoredResource(simpy.Resource):
...     def __init__(self, *args, **kwargs):
...         super().__init__(*args, **kwargs)
...         self.data = []
...
...     def request(self, *args, **kwargs):
...         self.data.append((self._env.now, len(self.queue)))
...         return super().request(*args, **kwargs)
...
...     def release(self, *args, **kwargs):
...         self.data.append((self._env.now, len(self.queue)))
...         return super().release(*args, **kwargs)
>>>
>>> def test_process(env, res):
...     with res.request() as req:
...         yield req
...         yield env.timeout(1)
>>>
>>> env = simpy.Environment()
>>>
>>> res = MonitoredResource(env, capacity=1)
>>> p1 = env.process(test_process(env, res))
>>> p2 = env.process(test_process(env, res))
>>> env.run()
>>>
>>> print(res.data)
[(0, 0), (0, 0), (1, 1), (2, 0)]
```

Ao contrário do primeiro exemplo, agora não se tem um só recurso modificado/patched mas sim uma classe toda. Também foi removida toda a parte considerada adaptável do primeiro exemplo: Só é monitorado o `Resource` nomeado como tal, só é coletado o dado *antes* das requisições serem feitas e só é coletado/registrado o momento/tempo e o tamanho da fila. Porém, você precisaria de menos da metade do código.

## Rastrear Eventos

Com o objetivo de fazer um debug ou visualizar uma simulação, você pode querer rastrear quando os eventos são criados, acionados e processados. Talvez você também queira rastrear quais processos criaram um determinado evento e que processos aguardaram um determinado evento.

As duas funções mais interessantes para esse caso de uso são `Environment.step()`, onde todos os eventos são processados e `Environment.schedule()`, onde todos os eventos são agendados e inseridos em uma fila de eventos do Simpy.

Aqui temos um exemplo que demonstra como `Environment.step()` pode ser adaptado de modo a rastrear todos os eventos processados:

```

>>> from functools import partial, wraps
>>> import simpy
>>>
>>> def trace(env, callback):
...     """Substitui o método ``step()`` de *env* por uma função de rastreamento
...     que executa *callbacks* informando o momento do evento, prioridade,
...     identificação e sua instância antes de que ela seja processada.
...
...     """
...     def get_wrapper(env_step, callback):
...         """Gerando o wrapper para o método env.step()."""
...         @wraps(env_step)
...         def tracing_step():
...             """Executa *callback* para o próximo evento, caso tenha
...             ocorrido um antes executando ``env.step()``."""
...             if len(env._queue):
...                 t, prio, eid, event = env._queue[0]
...                 callback(t, prio, eid, event)
...             return env_step()
...         return tracing_step
...
...     env.step = get_wrapper(env.step, callback)
>>>
>>> def monitor(data, t, prio, eid, event):
...     data.append((t, eid, type(event)))
>>>
>>> def test_process(env):
...     yield env.timeout(1)
>>>
>>> data = []
>>> # Vincular *data* como o primeiro argumento de monitor()
>>> # mais em https://docs.python.org/3/library/functools.html#functools.partial
>>> monitor = partial(monitor, data)
>>>
>>> env = simpy.Environment()
>>> trace(env, monitor)
>>>
>>> p = env.process(test_process(env))
>>> env.run(until=p)
>>>
>>> for d in data:
...     print(d)
(0, 0, <class 'simpy.events.Initialize'>)
(1, 1, <class 'simpy.events.Timeout'>)
(1, 2, <class 'simpy.events.Process'>)
    
```

O exemplo acima é baseado em um pull request de Steve Pothier.

Usando os mesmos conceitos, você também pode fazer um patch do método `Environment.schedule()`. Isso lhe dará acesso ao centro da informação quando determinado evento é agendado e para que instante.

Além disso, você também pode modificar algumas ou todas as classes de evento do SimPy, por ex. o método `__init__()` a fim de rastrear quando e como um evento está sendo criado.

## Time and Scheduling

The aim of this section is to give you a deeper understanding of how time passes in SimPy and how it schedules and processes events.

### What is time?

*Time* itself is not easy to grasp. The [wikipedians describe it](#) this way:

*«Time is the indefinite continued progress of existence and events that occur in apparently irreversible succession from the past through the present to the future. Time is a component quantity of various measurements used to sequence events, to compare the duration of events or the intervals between them, and to quantify rates of change of quantities in material reality or in the conscious experience. Time is often referred to as the fourth dimension, along with the three spatial dimensions.»*

### What's the problem with it?

Often, events (in the real world) appear to happen “at the same time”, when they are in fact happening at slightly different times. Here is an obvious example: Alice and Bob have birthday on the same day. If your time scale is in days, both birthday events happen at the same time. If you increase the resolution of your clock, e.g. to minutes, you may realise that Alice was actually born at 0:42 in the morning and Bob at 11:14 and that there's quite a difference between the time of both events.

Doing simulation on computers suffers from similar problems. Integers (and floats, too) are discrete numbers with a lot of void in between them. Thus, events that would occur after each other in the real world (e.g., at  $t_1 = 0.1$  and  $t_2 = 0.2$ ) might occur at the “same” time if mapped to an integer scale (e.g., at  $t = 0$ ).

On the other hand, SimPy is (like most simulation frameworks) a single-threaded, deterministic library. It processes events sequentially – one after another. If two events are scheduled at the same time, the one that is scheduled first will also be the processed first (FIFO).

That is very important for you to understand. The processes in your modeled/simulated world may run “in parallel”, but when the simulation runs on your CPU, all events are processed sequentially and deterministically. If you run your simulation multiple times (and if you don't use `random` ;-)), you will *always* get the same results.

So keep this in mind:

- In the real world, there's usually no *at the same time*.
- Discretization of the time scale can make events appear to be *at the same time*.
- SimPy processes events *one after another*, even if they have the *same time*.

### SimPy Events and time

Before we continue, let's recap the states an event can be in (see [Events](#) for details):

- untriggered: not known to the event queue
- triggered: scheduled at a time  $t$  and inserted into the event queue
- processed: removed from the event queue

SimPy's event queue is implemented as a [heap queue](#): “Heaps are binary trees for which every parent node has a value less than or equal to any of its children.” So if we insert events as tuples  $(t, event)$  (with  $t$  being the scheduled time) into it, the first element in the queue will by definition always be the one with the smallest  $t$  and the next one to be processed.

However, storing  $(t, event)$  tuples will not work if two events are scheduled at the same time because events are not comparable. To fix this, we also store a strictly increasing event ID with them:  $(t, eid, event)$ . That way, if two events get scheduled for the same time, the one scheduled first will always be processed first.

## Porting from SimPy 2 to 3

Porting from SimPy 2 to SimPy 3 is not overly complicated. A lot of changes merely comprise copy/paste.

This guide describes the conceptual and API changes between both SimPy versions and shows you how to change your code for SimPy 3.

### Imports

In SimPy 2, you had to decide at import-time whether you wanted to use a normal simulation (`SimPy.Simulation`), a real-time simulation (`SimPy.SimulationRT`) or something else. You usually had to import `Simulation` (or `SimulationRT`), `Process` and some of the SimPy keywords (`hold` or `passivate`, for example) from that package.

In SimPy 3, you usually need to import much less classes and modules (for example, all keywords are gone). In most use cases you will now only need to import `simpy`.

#### SimPy 2

```
from SimPy.Simulation import Simulation, Process, hold
```

#### SimPy 3

```
import simpy
```

### The `Simulation*` classes

SimPy 2 encapsulated the simulation state in a `Simulation*` class (e.g., `Simulation`, `SimulationRT` or `SimulationTrace`). This class also had a `simulate()` method that executed a normal simulation, a real-time simulation or something else (depending on the particular class).

There was a global `Simulation` instance that was automatically created when you imported SimPy. You could also instantiate it on your own to use SimPy's object-oriented API. This led to some confusion and problems, because you had to pass the `Simulation` instance around when you were using the object-oriented API but not if you were using the procedural API.

In SimPy 3, an `Environment` replaces `Simulation` and `RealtimeEnvironment` replaces `SimulationRT`. You always need to instantiate an environment. There's no more global state.

To execute a simulation, you call the environment's `run()` method.

#### SimPy 2

```
# Procedural API
from SimPy.Simulation import initialize, simulate

initialize()
# Start processes
simulate(until=10)
```

```
# Object-oriented API
from SimPy.Simulation import Simulation

sim = Simulation()
# Start processes
sim.simulate(until=10)
```

### SimPy3

```
import simpy

env = simpy.Environment()
# Start processes
env.run(until=10)
```

## Defining a Process

Processes had to inherit the `Process` base class in SimPy 2. Subclasses had to implement at least a so called *Process Execution Method (PEM)* (which is basically a generator function) and in most cases `__init__()`. Each process needed to know the `Simulation` instance it belonged to. This reference was passed implicitly in the procedural API and had to be passed explicitly in the object-oriented API. Apart from some internal problems, this made it quite cumbersome to define a simple process.

Processes were started by passing the `Process` and a generator instance created by the generator function to either the global `activate()` function or the corresponding `Simulation` method.

A process in SimPy 3 is a Python generator (no matter if it's defined on module level or as an instance method) wrapped in a `Process` instance. The generator usually requires a reference to a `Environment` to interact with, but this is completely optional.

Processes are can be started by creating a `Process` instance and passing the generator to it. The environment provides a shortcut for this: `process()`.

### SimPy 2

```
# Procedural API
from SimPy.Simulation import Process

class MyProcess(Process):
    def __init__(self, another_param):
        super().__init__()
        self.another_param = another_param

    def generator_function(self):
        """Implement the process' behavior."""
        yield something

initialize()
proc = Process('Spam')
activate(proc, proc.generator_function())
```

```
# Object-oriented API
from SimPy.Simulation import Simulation, Process

class MyProcess(Process):
    def __init__(self, sim, another_param):
        super().__init__(sim=sim)
        self.another_param = another_param
```

```

def generator_function(self):
    """Implement the process' behaviour."""
    yield something

sim = Simulation()
proc = Process(sim, 'Spam')
sim.activate(proc, proc.generator_function())

```

### SimPy 3

```

import simpy

def generator_function(env, another_param):
    """Implement the process' behavior."""
    yield something

env = simpy.Environment()
proc = env.process(generator_function(env, 'Spam'))

```

## SimPy Keywords (hold etc.)

In SimPy 2, processes created new events by yielding a *SimPy Keyword* and some additional parameters (at least self). These keywords had to be imported from `SimPy.Simulation*` if they were used. Internally, the keywords were mapped to a function that generated the according event.

In SimPy 3, you directly yield *events* if you want to wait for an event to occur. You can instantiate an event directly or use the shortcuts provided by *Environment*.

Generally, whenever a process yields an event, the execution of the process is suspended and resumed once the event has been triggered. To motivate this understanding, some of the events were renamed. For example, the `hold` keyword meant to wait until some time has passed. In terms of events this means that a timeout has happened. Therefore `hold` has been replaced by a *Timeout* event.

---

**Nota:** *Process* is also an *Event*. If you want to wait for a process to finish, simply yield it.

---

### SimPy 2

```

yield hold, self, duration
yield passivate, self
yield request, self, resource
yield release, self, resource
yield waitevent, self, event
yield waitevent, self, [event_a, event_b, event_c]
yield queueevent, self, event_list
yield get, self, level, amount
yield put, self, level, amount

```

### SimPy 3

```

yield env.timeout(duration)           # hold: renamed
yield env.event()                     # passivate: renamed
yield resource.request()              # Request is now bound to class Resource
resource.release()                    # Release no longer needs to be yielded
yield event                           # waitevent: just yield the event
yield env.all_of([event_a, event_b, event_c]) # waitvent
yield env.any_of([event_a, event_b, event_c]) # queueevent

```

```

yield container.get(amount)      # Level is now called Container
yield container.put(amount)

yield event_a | event_b         # Wait for either a or b. This is new.
yield event_a & event_b         # Wait for a and b. This is new.
yield env.process(calculation(env)) # Wait for the process calculation to
                                  # to finish.

```

## Partially supported features

The following `waituntil` keyword is not completely supported anymore:

```
yield waituntil, self, cond_func
```

SimPy 2 was evaluating `cond_func` after *every* event, which was computationally very expensive. One possible workaround is for example the following process, which evaluates `cond_func` periodically:

```

def waituntil(env, cond_func, delay=1):
    while not cond_func():
        yield env.timeout(delay)

# Usage:
yield waituntil(env, cond_func)

```

## Interrupts

In SimPy 2, `interrupt()` was a method of the interrupting process. The victim of the interrupt had to be passed as an argument.

The victim was not directly notified of the interrupt but had to check if the `interrupted` flag was set. Afterwards, it had to reset the interrupt via `interruptReset()`. You could manually set the `interruptCause` attribute of the victim.

Explicitly checking for an interrupt is obviously error prone as it is too easy to be forgotten.

In SimPy 3, you call `interrupt()` on the victim process. You can optionally supply a cause. An `Interrupt` is then thrown into the victim process, which has to handle the interrupt via `try: ... except Interrupt: ...`

### SimPy 2

```

class Interrupter(Process):
    def __init__(self, victim):
        super().__init__()
        self.victim = victim

    def run(self):
        yield hold, self, 1
        self.interrupt(self.victim_proc)
        self.victim_proc.interruptCause = 'Spam'

class Victim(Process):
    def run(self):
        yield hold, self, 10
        if self.interrupted:
            cause = self.interruptCause
            self.interruptReset()

```

### SimPy 3

```
def interrupter(env, victim_proc):  
    yield env.timeout(1)  
    victim_proc.interrupt('Spam')  
  
def victim(env):  
    try:  
        yield env.timeout(10)  
    except Interrupt as interrupt:  
        cause = interrupt.cause
```

### Conclusion

This guide is by no means complete. If you run into problems, please have a look at the other [guides](#), the [examples](#) or the [API Reference](#). You are also very welcome to submit improvements. Just create a pull request at [bitbucket](#).

---

**Examples**

---

All theory is grey. In this section, we present various practical examples that demonstrate how to use SimPy's features. Here's a list of examples grouped by features they demonstrate.

**Condition events**

- Bank Renege
- Movie Renege

**Interrupts**

- Machine Shop

**Monitoring****Resources: Container**

- Gas Station Refueling

**Resources: Preemptive Resource**

- Machine Shop

**Resources: Resource**

- Bank Renege
- Carwash
- Gas Station Refueling
- Movie Renege

## Resources: Store

- Event Latency
- Process Communication

## Shared events

- Movie Reneger

## Waiting for other processes

- Carwash
- Gas Station Refueling

## All examples

### Bank Reneger

Covers:

- Resources: Resource
- Condition events

A counter with a random service time and customers who renege. Based on the program `bank08.py` from TheBank tutorial of SimPy 2. (KGM)

This example models a bank counter and customers arriving t random times. Each customer has a certain patience. It waits to get to the counter until she's at the end of her tether. If she gets to the counter, she uses it for a while before releasing it.

New customers are created by the `source` process every few time steps.

```
"""
Bank renege example

Covers:
- Resources: Resource
- Condition events

Scenario:
A counter with a random service time and customers who renege. Based on the
program bank08.py from TheBank tutorial of SimPy 2. (KGM)

"""
import random

import simpy
```

```

RANDOM_SEED = 42
NEW_CUSTOMERS = 5 # Total number of customers
INTERVAL_CUSTOMERS = 10.0 # Generate new customers roughly every x seconds
MIN_PATIENCE = 1 # Min. customer patience
MAX_PATIENCE = 3 # Max. customer patience

def source(env, number, interval, counter):
    """Source generates customers randomly"""
    for i in range(number):
        c = customer(env, 'Customer%02d' % i, counter, time_in_bank=12.0)
        env.process(c)
        t = random.expovariate(1.0 / interval)
        yield env.timeout(t)

def customer(env, name, counter, time_in_bank):
    """Customer arrives, is served and leaves."""
    arrive = env.now
    print('%7.4f %s: Here I am' % (arrive, name))

    with counter.request() as req:
        patience = random.uniform(MIN_PATIENCE, MAX_PATIENCE)
        # Wait for the counter or abort at the end of our tether
        results = yield req | env.timeout(patience)

        wait = env.now - arrive

        if req in results:
            # We got to the counter
            print('%7.4f %s: Waited %6.3f' % (env.now, name, wait))

            tib = random.expovariate(1.0 / time_in_bank)
            yield env.timeout(tib)
            print('%7.4f %s: Finished' % (env.now, name))

        else:
            # We reneged
            print('%7.4f %s: RENEGED after %6.3f' % (env.now, name, wait))

# Setup and start the simulation
print('Bank renege')
random.seed(RANDOM_SEED)
env = simpy.Environment()

# Start processes and run
counter = simpy.Resource(env, capacity=1)
env.process(source(env, NEW_CUSTOMERS, INTERVAL_CUSTOMERS, counter))
env.run()

```

The simulation's output:

```

Bank renege
0.0000 Customer00: Here I am
0.0000 Customer00: Waited 0.000
3.8595 Customer00: Finished
10.2006 Customer01: Here I am
10.2006 Customer01: Waited 0.000

```

```

12.7265 Customer02: Here I am
13.9003 Customer02: RENEGED after 1.174
23.7507 Customer01: Finished
34.9993 Customer03: Here I am
34.9993 Customer03: Waited 0.000
37.9599 Customer03: Finished
40.4798 Customer04: Here I am
40.4798 Customer04: Waited 0.000
43.1401 Customer04: Finished

```

## Carwash

Covers:

- Waiting for other processes
- Resources: Resource

The *Carwash* example is a simulation of a carwash with a limited number of machines and a number of cars that arrive at the carwash to get cleaned.

The carwash uses a *Resource* to model the limited number of washing machines. It also defines a process for washing a car.

When a car arrives at the carwash, it requests a machine. Once it got one, it starts the carwash's *wash* processes and waits for it to finish. It finally releases the machine and leaves.

The cars are generated by a *setup* process. After creating an initial amount of cars it creates new *car* processes after a random time interval as long as the simulation continues.

```

"""
Carwash example.

Covers:

- Waiting for other processes
- Resources: Resource

Scenario:
    A carwash has a limited number of washing machines and defines
    a washing processes that takes some (random) time.

    Car processes arrive at the carwash at a random time. If one washing
    machine is available, they start the washing process and wait for it
    to finish. If not, they wait until they an use one.

"""
import random

import simpy

RANDOM_SEED = 42
NUM_MACHINES = 2 # Number of machines in the carwash
WASHTIME = 5    # Minutes it takes to clean a car
T_INTER = 7     # Create a car every ~7 minutes
SIM_TIME = 20   # Simulation time in minutes

```

```

class Carwash(object):
    """A carwash has a limited number of machines (`NUM_MACHINES`) to
    clean cars in parallel.

    Cars have to request one of the machines. When they got one, they
    can start the washing processes and wait for it to finish (which
    takes `washtime` minutes).

    """
    def __init__(self, env, num_machines, washtime):
        self.env = env
        self.machine = simpy.Resource(env, num_machines)
        self.washtime = washtime

    def wash(self, car):
        """The washing processes. It takes a `car` processes and tries
        to clean it."""
        yield self.env.timeout(WASHTIME)
        print("Carwash removed %d%% of %s's dirt." %
              (random.randint(50, 99), car))

def car(env, name, cw):
    """The car process (each car has a `name`) arrives at the carwash
    (`cw`) and requests a cleaning machine.

    It then starts the washing process, waits for it to finish and
    leaves to never come back ...

    """
    print('%s arrives at the carwash at %.2f.' % (name, env.now))
    with cw.machine.request() as request:
        yield request

        print('%s enters the carwash at %.2f.' % (name, env.now))
        yield env.process(cw.wash(name))

        print('%s leaves the carwash at %.2f.' % (name, env.now))

def setup(env, num_machines, washtime, t_inter):
    """Create a carwash, a number of initial cars and keep creating cars
    approx. every `t_inter` minutes."""
    # Create the carwash
    carwash = Carwash(env, num_machines, washtime)

    # Create 4 initial cars
    for i in range(4):
        env.process(car(env, 'Car %d' % i, carwash))

    # Create more cars while the simulation is running
    while True:
        yield env.timeout(random.randint(t_inter - 2, t_inter + 2))
        i += 1
        env.process(car(env, 'Car %d' % i, carwash))

# Setup and start the simulation

```

```

print('Carwash')
print('Check out http://youtu.be/fXXmeP9TvBg while simulating ... ;-)')
random.seed(RANDOM_SEED) # This helps reproducing the results

# Create an environment and start the setup process
env = simpy.Environment()
env.process(setup(env, NUM_MACHINES, WASHTIME, T_INTER))

# Execute!
env.run(until=SIM_TIME)

```

The simulation's output:

```

Carwash
Check out http://youtu.be/fXXmeP9TvBg while simulating ... ;-)
Car 0 arrives at the carwash at 0.00.
Car 1 arrives at the carwash at 0.00.
Car 2 arrives at the carwash at 0.00.
Car 3 arrives at the carwash at 0.00.
Car 0 enters the carwash at 0.00.
Car 1 enters the carwash at 0.00.
Car 4 arrives at the carwash at 5.00.
Carwash removed 97% of Car 0's dirt.
Carwash removed 67% of Car 1's dirt.
Car 0 leaves the carwash at 5.00.
Car 1 leaves the carwash at 5.00.
Car 2 enters the carwash at 5.00.
Car 3 enters the carwash at 5.00.
Car 5 arrives at the carwash at 10.00.
Carwash removed 64% of Car 2's dirt.
Carwash removed 58% of Car 3's dirt.
Car 2 leaves the carwash at 10.00.
Car 3 leaves the carwash at 10.00.
Car 4 enters the carwash at 10.00.
Car 5 enters the carwash at 10.00.
Carwash removed 97% of Car 4's dirt.
Carwash removed 56% of Car 5's dirt.
Car 4 leaves the carwash at 15.00.
Car 5 leaves the carwash at 15.00.
Car 6 arrives at the carwash at 16.00.
Car 6 enters the carwash at 16.00.

```

## Machine Shop

Covers:

- Interrupts
- Resources: PreemptiveResource

This example comprises a workshop with  $n$  identical machines. A stream of jobs (enough to keep the machines busy) arrives. Each machine breaks down periodically. Repairs are carried out by one repairman. The repairman has other, less important tasks to perform, too. Broken machines preempt these tasks. The repairman continues them when he is done with the machine repair. The workshop works continuously.

A machine has two processes: *working* implements the actual behaviour of the machine (producing parts). *break\_machine* periodically interrupts the *working* process to simulate the machine failure.

The repairman's other job is also a process (implemented by *other\_job*). The repairman itself is a *PreemptiveResource* with a capacity of 1. The machine repairing has a priority of 1, while the other job has a priority of 2 (the smaller the number, the higher the priority).

```

"""
Machine shop example

Covers:

- Interrupts
- Resources: PreemptiveResource

Scenario:
  A workshop has *n* identical machines. A stream of jobs (enough to
  keep the machines busy) arrives. Each machine breaks down
  periodically. Repairs are carried out by one repairman. The repairman
  has other, less important tasks to perform, too. Broken machines
  preempt these tasks. The repairman continues them when he is done
  with the machine repair. The workshop works continuously.

"""
import random

import simpy

RANDOM_SEED = 42
PT_MEAN = 10.0          # Avg. processing time in minutes
PT_SIGMA = 2.0         # Sigma of processing time
MTTF = 300.0           # Mean time to failure in minutes
BREAK_MEAN = 1 / MTTF # Param. for expovariate distribution
REPAIR_TIME = 30.0     # Time it takes to repair a machine in minutes
JOB_DURATION = 30.0    # Duration of other jobs in minutes
NUM_MACHINES = 10     # Number of machines in the machine shop
WEEKS = 4              # Simulation time in weeks
SIM_TIME = WEEKS * 7 * 24 * 60 # Simulation time in minutes

def time_per_part():
    """Return actual processing time for a concrete part."""
    return random.normalvariate(PT_MEAN, PT_SIGMA)

def time_to_failure():
    """Return time until next failure for a machine."""
    return random.expovariate(BREAK_MEAN)

class Machine(object):
    """A machine produces parts and my get broken every now and then.

    If it breaks, it requests a *repairman* and continues the production
    after the it is repaired.

    A machine has a *name* and a numberof *parts_made* thus far.

    """
    def __init__(self, env, name, repairman):
        self.env = env

```

```

self.name = name
self.parts_made = 0
self.broken = False

# Start "working" and "break_machine" processes for this machine.
self.process = env.process(self.working(repairman))
env.process(self.break_machine())

def working(self, repairman):
    """Produce parts as long as the simulation runs.

    While making a part, the machine may break multiple times.
    Request a repairman when this happens.

    """
    while True:
        # Start making a new part
        done_in = time_per_part()
        while done_in:
            try:
                # Working on the part
                start = self.env.now
                yield self.env.timeout(done_in)
                done_in = 0 # Set to 0 to exit while loop.

            except simpy.Interrupt:
                self.broken = True
                done_in -= self.env.now - start # How much time left?

                # Request a repairman. This will preempt its "other_job".
                with repairman.request(priority=1) as req:
                    yield req
                    yield self.env.timeout(REPAIR_TIME)

                self.broken = False

        # Part is done.
        self.parts_made += 1

def break_machine(self):
    """Break the machine every now and then."""
    while True:
        yield self.env.timeout(time_to_failure())
        if not self.broken:
            # Only break the machine if it is currently working.
            self.process.interrupt()

def other_jobs(env, repairman):
    """The repairman's other (unimportant) job."""
    while True:
        # Start a new job
        done_in = JOB_DURATION
        while done_in:
            # Retry the job until it is done.
            # It's priority is lower than that of machine repairs.
            with repairman.request(priority=2) as req:
                yield req

```

```

    try:
        start = env.now
        yield env.timeout(done_in)
        done_in = 0
    except simpy.Interrupt:
        done_in -= env.now - start

# Setup and start the simulation
print('Machine shop')
random.seed(RANDOM_SEED) # This helps reproducing the results

# Create an environment and start the setup process
env = simpy.Environment()
repairman = simpy.PreemptiveResource(env, capacity=1)
machines = [Machine(env, 'Machine %d' % i, repairman)
             for i in range(NUM_MACHINES)]
env.process(other_jobs(env, repairman))

# Execute!
env.run(until=SIM_TIME)

# Analysis/results
print('Machine shop results after %s weeks' % WEEKS)
for machine in machines:
    print('%s made %d parts.' % (machine.name, machine.parts_made))

```

The simulation's output:

```

Machine shop
Machine shop results after 4 weeks
Machine 0 made 3251 parts.
Machine 1 made 3273 parts.
Machine 2 made 3242 parts.
Machine 3 made 3343 parts.
Machine 4 made 3387 parts.
Machine 5 made 3244 parts.
Machine 6 made 3269 parts.
Machine 7 made 3185 parts.
Machine 8 made 3302 parts.
Machine 9 made 3279 parts.

```

## Movie Renege

Covers:

- Resources: Resource
- Condition events
- Shared events

This examples models a movie theater with one ticket counter selling tickets for three movies (next show only). People arrive at random times and try to buy a random number (1–6) tickets for a random movie. When a movie is sold out, all people waiting to buy a ticket for that movie renege (leave the queue).

The movie theater is just a container for all the related data (movies, the counter, tickets left, collected data, ...). The counter is a *Resource* with a capacity of one.

The *moviegoer* process starts waiting until either it's his turn (it acquires the counter resource) or until the *sold out* signal is triggered. If the latter is the case it reneges (leaves the queue). If it gets to the counter, it tries to buy some tickets. This might not be successful, e.g. if the process tries to buy 5 tickets but only 3 are left. If less than two tickets are left after the ticket purchase, the *sold out* signal is triggered.

Moviegoers are generated by the *customer arrivals* process. It also chooses a movie and the number of tickets for the moviegoer.

```

"""
Movie renege example

Covers:

- Resources: Resource
- Condition events
- Shared events

Scenario:
  A movie theatre has one ticket counter selling tickets for three
  movies (next show only). When a movie is sold out, all people waiting
  to buy tickets for that movie renege (leave queue).

"""
import collections
import random

import simpy

RANDOM_SEED = 42
TICKETS = 50 # Number of tickets per movie
SIM_TIME = 120 # Simulate until

def moviegoer(env, movie, num_tickets, theater):
    """A moviegoer tries to buy a number of tickets (*num_tickets*) for
    a certain *movie* in a *theater*.


If the movie becomes sold out, she leaves the theater. If she gets
    to the counter, she tries to buy a number of tickets. If not enough
    tickets are left, she argues with the teller and leaves.



If at most one ticket is left after the moviegoer bought her
    tickets, the *sold out* event for this movie is triggered causing
    all remaining moviegoers to leave.


    """
    with theater.counter.request() as my_turn:
        # Wait until its our turn or until the movie is sold out
        result = yield my_turn | theater.sold_out[movie]

        # Check if it's our turn or if movie is sold out
        if my_turn not in result:
            theater.num_renegers[movie] += 1
            env.exit()

        # Check if enough tickets left.
        if theater.available[movie] < num_tickets:
            # Moviegoer leaves after some discussion

```

```

        yield env.timeout(0.5)
        env.exit()

    # Buy tickets
    theater.available[movie] -= num_tickets
    if theater.available[movie] < 2:
        # Trigger the "sold out" event for the movie
        theater.sold_out[movie].succeed()
        theater.when_sold_out[movie] = env.now
        theater.available[movie] = 0
    yield env.timeout(1)

def customer_arrivals(env, theater):
    """Create new *moviegoers* until the sim time reaches 120."""
    while True:
        yield env.timeout(random.expovariate(1 / 0.5))

        movie = random.choice(theater.movies)
        num_tickets = random.randint(1, 6)
        if theater.available[movie]:
            env.process(moviegoer(env, movie, num_tickets, theater))

Theater = collections.namedtuple('Theater', 'counter, movies, available, '
                                'sold_out, when_sold_out, '
                                'num_renegers')

# Setup and start the simulation
print('Movie renege')
random.seed(RANDOM_SEED)
env = simpy.Environment()

# Create movie theater
counter = simpy.Resource(env, capacity=1)
movies = ['Python Unchained', 'Kill Process', 'Pulp Implementation']
available = {movie: TICKETS for movie in movies}
sold_out = {movie: env.event() for movie in movies}
when_sold_out = {movie: None for movie in movies}
num_renegers = {movie: 0 for movie in movies}
theater = Theater(counter, movies, available, sold_out, when_sold_out,
                 num_renegers)

# Start process and run
env.process(customer_arrivals(env, theater))
env.run(until=SIM_TIME)

# Analysis/results
for movie in movies:
    if theater.sold_out[movie]:
        print('Movie "%s" sold out %.1f minutes after ticket counter '
              'opening.' % (movie, theater.when_sold_out[movie]))
        print('  Number of people leaving queue when film sold out: %s' %
              theater.num_renegers[movie])

```

The simulation's output:

```

Movie renege
Movie "Python Unchained" sold out 38.0 minutes after ticket counter opening.
  Number of people leaving queue when film sold out: 16
Movie "Kill Process" sold out 43.0 minutes after ticket counter opening.
  Number of people leaving queue when film sold out: 5
Movie "Pulp Implementation" sold out 28.0 minutes after ticket counter opening.
  Number of people leaving queue when film sold out: 5

```

## Gas Station Refueling

Covers:

- Resources: Resource
- Resources: Container
- Waiting for other processes

This examples models a gas station and cars that arrive at the station for refueling.

The gas station has a limited number of fuel pumps and a fuel tank that is shared between the fuel pumps. The gas station is thus modeled as *Resource*. The shared fuel tank is modeled with a *Container*.

Vehicles arriving at the gas station first request a fuel pump from the station. Once they acquire one, they try to take the desired amount of fuel from the fuel pump. They leave when they are done.

The gas stations fuel level is regularly monitored by *gas station control*. When the level drops below a certain threshold, a *tank truck* is called to refuel the gas station itself.

```

"""
Gas Station Refueling example

Covers:

- Resources: Resource
- Resources: Container
- Waiting for other processes

Scenario:
  A gas station has a limited number of gas pumps that share a common
  fuel reservoir. Cars randomly arrive at the gas station, request one
  of the fuel pumps and start refueling from that reservoir.

  A gas station control process observes the gas station's fuel level
  and calls a tank truck for refueling if the station's level drops
  below a threshold.

"""
import itertools
import random

import simpy

RANDOM_SEED = 42
GAS_STATION_SIZE = 200      # liters
THRESHOLD = 10             # Threshold for calling the tank truck (in %)
FUEL_TANK_SIZE = 50        # liters
FUEL_TANK_LEVEL = [5, 25] # Min/max levels of fuel tanks (in liters)

```

```

REFUELING_SPEED = 2          # liters / second
TANK_TRUCK_TIME = 300       # Seconds it takes the tank truck to arrive
T_INTER = [30, 300]        # Create a car every [min, max] seconds
SIM_TIME = 1000            # Simulation time in seconds

def car(name, env, gas_station, fuel_pump):
    """A car arrives at the gas station for refueling.

    It requests one of the gas station's fuel pumps and tries to get the
    desired amount of gas from it. If the stations reservoir is
    depleted, the car has to wait for the tank truck to arrive.

    """
    fuel_tank_level = random.randint(*FUEL_TANK_LEVEL)
    print('%s arriving at gas station at %.1f' % (name, env.now))
    with gas_station.request() as req:
        start = env.now
        # Request one of the gas pumps
        yield req

        # Get the required amount of fuel
        liters_required = FUEL_TANK_SIZE - fuel_tank_level
        yield fuel_pump.get(liters_required)

        # The "actual" refueling process takes some time
        yield env.timeout(liters_required / REFUELING_SPEED)

    print('%s finished refueling in %.1f seconds.' % (name,
                                                    env.now - start))

def gas_station_control(env, fuel_pump):
    """Periodically check the level of the *fuel_pump* and call the tank
    truck if the level falls below a threshold."""
    while True:
        if fuel_pump.level / fuel_pump.capacity * 100 < THRESHOLD:
            # We need to call the tank truck now!
            print('Calling tank truck at %d' % env.now)
            # Wait for the tank truck to arrive and refuel the station
            yield env.process(tank_truck(env, fuel_pump))

        yield env.timeout(10) # Check every 10 seconds

def tank_truck(env, fuel_pump):
    """Arrives at the gas station after a certain delay and refuels it."""
    yield env.timeout(TANK_TRUCK_TIME)
    print('Tank truck arriving at time %d' % env.now)
    ammount = fuel_pump.capacity - fuel_pump.level
    print('Tank truck refuelling %.1f liters.' % ammount)
    yield fuel_pump.put(ammount)

def car_generator(env, gas_station, fuel_pump):
    """Generate new cars that arrive at the gas station."""
    for i in itertools.count():
        yield env.timeout(random.randint(*T_INTER))

```

```

        env.process(car('Car %d' % i, env, gas_station, fuel_pump))

# Setup and start the simulation
print('Gas Station refuelling')
random.seed(RANDOM_SEED)

# Create environment and start processes
env = simpy.Environment()
gas_station = simpy.Resource(env, 2)
fuel_pump = simpy.Container(env, GAS_STATION_SIZE, init=GAS_STATION_SIZE)
env.process(gas_station_control(env, fuel_pump))
env.process(car_generator(env, gas_station, fuel_pump))

# Execute!
env.run(until=SIM_TIME)

```

The simulation's output:

```

Gas Station refuelling
Car 0 arriving at gas station at 87.0
Car 0 finished refueling in 18.5 seconds.
Car 1 arriving at gas station at 129.0
Car 1 finished refueling in 19.0 seconds.
Car 2 arriving at gas station at 284.0
Car 2 finished refueling in 21.0 seconds.
Car 3 arriving at gas station at 385.0
Car 3 finished refueling in 13.5 seconds.
Car 4 arriving at gas station at 459.0
Calling tank truck at 460
Car 4 finished refueling in 22.0 seconds.
Car 5 arriving at gas station at 705.0
Car 6 arriving at gas station at 750.0
Tank truck arriving at time 760
Tank truck refuelling 188.0 liters.
Car 6 finished refueling in 29.0 seconds.
Car 5 finished refueling in 76.5 seconds.
Car 7 arriving at gas station at 891.0
Car 7 finished refueling in 13.0 seconds.

```

## Process Communication

Covers:

- Resources: Store

This example shows how to interconnect simulation model elements together using “resources.Store” for one-to-one, and many-to-one asynchronous processes. For one-to-many a simple BroadCastPipe class is constructed from Store.

**When Useful:** When a consumer process does not always wait on a generating process and these processes run asynchronously. This example shows how to create a buffer and also tell is the consumer process was late yielding to the event from a generating process.

This is also useful when some information needs to be broadcast to many receiving processes

Finally, using pipes can simplify how processes are interconnected to each other in a simulation model.

**Example By:** Keith Smith

```

"""
Process communication example

Covers:

- Resources: Store

Scenario:
    This example shows how to interconnect simulation model elements
    together using :class:`~simpy.resources.store.Store` for one-to-one,
    and many-to-one asynchronous processes. For one-to-many a simple
    BroadCastPipe class is constructed from Store.

When Useful:
    When a consumer process does not always wait on a generating process
    and these processes run asynchronously. This example shows how to
    create a buffer and also tell is the consumer process was late
    yielding to the event from a generating process.

    This is also useful when some information needs to be broadcast to
    many receiving processes

    Finally, using pipes can simplify how processes are interconnected to
    each other in a simulation model.

Example By:
    Keith Smith

"""
import random

import simpy

RANDOM_SEED = 42
SIM_TIME = 100

class BroadcastPipe(object):
    """A Broadcast pipe that allows one process to send messages to many.

    This construct is useful when message consumers are running at
    different rates than message generators and provides an event
    buffering to the consuming processes.

    The parameters are used to create a new
    :class:`~simpy.resources.store.Store` instance each time
    :meth:`~get_output_conn()` is called.

    """
    def __init__(self, env, capacity=simpy.core.Infinity):
        self.env = env
        self.capacity = capacity
        self.pipes = []

    def put(self, value):
        """Broadcast a *value* to all receivers."""
        if not self.pipes:

```

```

        raise RuntimeError('There are no output pipes.')
    events = [store.put(value) for store in self.pipes]
    return self.env.all_of(events) # Condition event for all "events"

def get_output_conn(self):
    """Get a new output connection for this broadcast pipe.

    The return value is a :class:`~simpy.resources.store.Store`.

    """
    pipe = simpy.Store(self.env, capacity=self.capacity)
    self.pipes.append(pipe)
    return pipe

def message_generator(name, env, out_pipe):
    """A process which randomly generates messages."""
    while True:
        # wait for next transmission
        yield env.timeout(random.randint(6, 10))

        # messages are time stamped to later check if the consumer was
        # late getting them. Note, using event.triggered to do this may
        # result in failure due to FIFO nature of simulation yields.
        # (i.e. if at the same env.now, message_generator puts a message
        # in the pipe first and then message_consumer gets from pipe,
        # the event.triggered will be True in the other order it will be
        # False
        msg = (env.now, '%s says hello at %d' % (name, env.now))
        out_pipe.put(msg)

def message_consumer(name, env, in_pipe):
    """A process which consumes messages."""
    while True:
        # Get event for message pipe
        msg = yield in_pipe.get()

        if msg[0] < env.now:
            # if message was already put into pipe, then
            # message_consumer was late getting to it. Depending on what
            # is being modeled this, may, or may not have some
            # significance
            print('LATE Getting Message: at time %d: %s received message: %s' %
                  (env.now, name, msg[1]))

        else:
            # message_consumer is synchronized with message_generator
            print('at time %d: %s received message: %s.' %
                  (env.now, name, msg[1]))

        # Process does some other work, which may result in missing messages
        yield env.timeout(random.randint(4, 8))

# Setup and start the simulation
print('Process communication')
random.seed(RANDOM_SEED)

```

```

env = simpy.Environment()

# For one-to-one or many-to-one type pipes, use Store
pipe = simpy.Store(env)
env.process(message_generator('Generator A', env, pipe))
env.process(message_consumer('Consumer A', env, pipe))

print('\nOne-to-one pipe communication\n')
env.run(until=SIM_TIME)

# For one-to many use BroadcastPipe
# (Note: could also be used for one-to-one, many-to-one or many-to-many)
env = simpy.Environment()
bc_pipe = BroadcastPipe(env)

env.process(message_generator('Generator A', env, bc_pipe))
env.process(message_consumer('Consumer A', env, bc_pipe.get_output_conn()))
env.process(message_consumer('Consumer B', env, bc_pipe.get_output_conn()))

print('\nOne-to-many pipe communication\n')
env.run(until=SIM_TIME)

```

The simulation's output:

```

Process communication

One-to-one pipe communication

at time 6: Consumer A received message: Generator A says hello at 6.
at time 12: Consumer A received message: Generator A says hello at 12.
at time 19: Consumer A received message: Generator A says hello at 19.
at time 26: Consumer A received message: Generator A says hello at 26.
at time 36: Consumer A received message: Generator A says hello at 36.
at time 46: Consumer A received message: Generator A says hello at 46.
at time 52: Consumer A received message: Generator A says hello at 52.
at time 58: Consumer A received message: Generator A says hello at 58.
LATE Getting Message: at time 66: Consumer A received message: Generator A says hello at 65
at time 75: Consumer A received message: Generator A says hello at 75.
at time 85: Consumer A received message: Generator A says hello at 85.
at time 95: Consumer A received message: Generator A says hello at 95.

One-to-many pipe communication

at time 10: Consumer A received message: Generator A says hello at 10.
at time 10: Consumer B received message: Generator A says hello at 10.
at time 18: Consumer A received message: Generator A says hello at 18.
at time 18: Consumer B received message: Generator A says hello at 18.
at time 27: Consumer A received message: Generator A says hello at 27.
at time 27: Consumer B received message: Generator A says hello at 27.
at time 34: Consumer A received message: Generator A says hello at 34.
at time 34: Consumer B received message: Generator A says hello at 34.
at time 40: Consumer A received message: Generator A says hello at 40.
LATE Getting Message: at time 41: Consumer B received message: Generator A says hello at 40
at time 46: Consumer A received message: Generator A says hello at 46.
LATE Getting Message: at time 47: Consumer B received message: Generator A says hello at 46
at time 56: Consumer A received message: Generator A says hello at 56.
at time 56: Consumer B received message: Generator A says hello at 56.
at time 65: Consumer A received message: Generator A says hello at 65.

```

```

at time 65: Consumer B received message: Generator A says hello at 65.
at time 74: Consumer A received message: Generator A says hello at 74.
at time 74: Consumer B received message: Generator A says hello at 74.
at time 82: Consumer A received message: Generator A says hello at 82.
at time 82: Consumer B received message: Generator A says hello at 82.
at time 92: Consumer A received message: Generator A says hello at 92.
at time 92: Consumer B received message: Generator A says hello at 92.
at time 98: Consumer B received message: Generator A says hello at 98.
at time 98: Consumer A received message: Generator A says hello at 98.

```

## Event Latency

Covers:

- Resources: Store

This example shows how to separate the time delay of events between processes from the processes themselves.

**When Useful:** When modeling physical things such as cables, RF propagation, etc. it better encapsulation to keep this propagation mechanism outside of the sending and receiving processes.

Can also be used to interconnect processes sending messages

**Example by:** Keith Smith

```

"""
Event Latency example

Covers:

- Resources: Store

Scenario:
    This example shows how to separate the time delay of events between
    processes from the processes themselves.

When Useful:
    When modeling physical things such as cables, RF propagation, etc. it
    better encapsulation to keep this propagation mechanism outside of the
    sending and receiving processes.

    Can also be used to interconnect processes sending messages

Example by:
    Keith Smith

"""
import simpy

SIM_DURATION = 100

class Cable(object):
    """This class represents the propagation through a cable."""
    def __init__(self, env, delay):
        self.env = env
        self.delay = delay
        self.store = simpy.Store(env)

```

```

def latency(self, value):
    yield self.env.timeout(self.delay)
    self.store.put(value)

def put(self, value):
    self.env.process(self.latency(value))

def get(self):
    return self.store.get()

def sender(env, cable):
    """A process which randomly generates messages."""
    while True:
        # wait for next transmission
        yield env.timeout(5)
        cable.put('Sender sent this at %d' % env.now)

def receiver(env, cable):
    """A process which consumes messages."""
    while True:
        # Get event for message pipe
        msg = yield cable.get()
        print('Received this at %d while %s' % (env.now, msg))

# Setup and start the simulation
print('Event Latency')
env = simpy.Environment()

cable = Cable(env, 10)
env.process(sender(env, cable))
env.process(receiver(env, cable))

env.run(until=SIM_DURATION)

```

The simulation's output:

```

Event Latency
Received this at 15 while Sender sent this at 5
Received this at 20 while Sender sent this at 10
Received this at 25 while Sender sent this at 15
Received this at 30 while Sender sent this at 20
Received this at 35 while Sender sent this at 25
Received this at 40 while Sender sent this at 30
Received this at 45 while Sender sent this at 35
Received this at 50 while Sender sent this at 40
Received this at 55 while Sender sent this at 45
Received this at 60 while Sender sent this at 50
Received this at 65 while Sender sent this at 55
Received this at 70 while Sender sent this at 60
Received this at 75 while Sender sent this at 65
Received this at 80 while Sender sent this at 70
Received this at 85 while Sender sent this at 75
Received this at 90 while Sender sent this at 80
Received this at 95 while Sender sent this at 85

```

You have ideas for better examples? Please send them to our [mailing list](#) or make a pull request on [bitbucket](#).

---

## API Reference

---

The API reference provides detailed descriptions of SimPy's classes and functions. It should be helpful if you plan to extend SimPy with custom components.

### simpy

The `simpy` module aggregates SimPy's most used components into a single namespace. This is purely for convenience. You can of course also access everything (and more!) via their actual submodules.

The following tables list all of the available components in this module.

#### Environments

<code>Environment([initial_time])</code>	Execution environment for an event-based simulation.
<code>RealtimeEnvironment([initial_time, factor, ...])</code>	Execution environment for an event-based simulation which is synchronized v

#### Events

<code>Event(env)</code>	An event that may happen at some point in time.
<code>Timeout(env, delay[, value])</code>	A <code>Event</code> that gets triggered after a <code>delay</code> has passed.
<code>Process(env, generator)</code>	Process an event yielding generator.
<code>AllOf(env, events)</code>	A <code>Condition</code> event that is triggered if all of a list of <code>events</code> have been successfully triggered.
<code>AnyOf(env, events)</code>	A <code>Condition</code> event that is triggered if any of a list of <code>events</code> has been successfully triggered.
<code>Interrupt(cause)</code>	Exception thrown into a process if it is interrupted (see <code>interrupt ()</code> ).

#### Resources

<code>Resource(env[, capacity])</code>	Resource with <code>capacity</code> of usage slots that can be requested by processes.
<code>PriorityResource(env[, capacity])</code>	A <code>Resource</code> supporting prioritized requests.
<code>PreemptiveResource(env[, capacity])</code>	A <code>PriorityResource</code> with preemption.
<code>Container(env[, capacity, init])</code>	Resource containing up to <code>capacity</code> of matter which may either be continuous (like water)
<code>Store(env[, capacity])</code>	Resource with <code>capacity</code> slots for storing arbitrary objects.
<code>PriorityItem</code>	Wrap an arbitrary <code>item</code> with an orderable <code>priority</code> .

Continu

Table 5.3 – continuação da página anterior

<code>PriorityStore(env[, capacity])</code>	Resource with <i>capacity</i> slots for storing objects in priority order.
<code>FilterStore(env[, capacity])</code>	Resource with <i>capacity</i> slots for storing arbitrary objects supporting filtered get requests.

## Exceptions

<code>SimPyException</code>	Base class for all SimPy specific exceptions.
<code>Interrupt(cause)</code>	Exception thrown into a process if it is interrupted (see <code>interrupt()</code> ).
<code>StopProcess(value)</code>	Raised to stop a SimPy process (similar to <code>StopIteration</code> ).

## simpy.core — SimPy’s core components

Core components for event-discrete simulation environments.

**class** `simpy.core.BaseEnvironment`

Base class for event processing environments.

An implementation must at least provide the means to access the current time of the environment (see `now`) and to schedule (see `schedule()`) events as well as processing them (see `step()`).

The class is meant to be subclassed for different execution environments. For example, SimPy defines a `Environment` for simulations with a virtual time and a `RealtimeEnvironment` that schedules and executes events in real (e.g., wallclock) time.

**now**

The current time of the environment.

**active\_process**

The currently active process of the environment.

**schedule** (*event*, *priority=1*, *delay=0*)

Schedule an *event* with a given *priority* and a *delay*.

There are two default priority values, `URGENT` and `NORMAL`.

**step()**

Processes the next event.

**run** (*until=None*)

Executes `step()` until the given criterion *until* is met.

- If it is `None` (which is the default), this method will return when there are no further events to be processed.
- If it is an `Event`, the method will continue stepping until this event has been triggered and will return its value. Raises a `RuntimeError` if there are no further events to be processed and the *until* event was not triggered.
- If it is a number, the method will continue stepping until the environment’s time reaches *until*.

**exit** (*value=None*)

Stop the current process, optionally providing a *value*.

This is a convenience function provided for Python versions prior to 3.3. From Python 3.3, you can instead use `return value` in a process.

---

**class** `simpy.core.Environment` (*initial\_time=0*)  
 Execution environment for an event-based simulation. The passing of time is simulated by stepping from event to event.

You can provide an *initial\_time* for the environment. By default, it starts at 0.

This class also provides aliases for common event types, for example *process*, *timeout* and *event*.

**now**  
 The current simulation time.

**active\_process**  
 The currently active process of the environment.

**process** (*generator*)  
 Create a new *Process* instance for *generator*.

**timeout** (*delay, value=None*)  
 Return a new *Timeout* event with a *delay* and, optionally, a *value*.

**event** ()  
 Return a new *Event* instance. Yielding this event suspends a process until another process triggers the event.

**all\_of** (*events*)  
 Return a new *AllOf* condition for a list of *events*.

**any\_of** (*events*)  
 Return a new *AnyOf* condition for a list of *events*.

**exit** (*value=None*)  
 Stop the current process, optionally providing a *value*.

This is a convenience function provided for Python versions prior to 3.3. From Python 3.3, you can instead use `return value` in a process.

**schedule** (*event, priority=1, delay=0*)  
 Schedule an *event* with a given *priority* and a *delay*.

**peek** ()  
 Get the time of the next scheduled event. Return *Infinity* if there is no further event.

**step** ()  
 Process the next event.

Raise an *EmptySchedule* if no further events are available.

**run** (*until=None*)  
 Executes *step()* until the given criterion *until* is met.

- If it is *None* (which is the default), this method will return when there are no further events to be processed.
- If it is an *Event*, the method will continue stepping until this event has been triggered and will return its value. Raises a *RuntimeError* if there are no further events to be processed and the *until* event was not triggered.
- If it is a number, the method will continue stepping until the environment's time reaches *until*.

**class** `simpy.core.BoundClass` (*cls*)  
 Allows classes to behave like methods.

The `__get__()` descriptor is basically identical to `function.__get__()` and binds the first argument of the *cls* to the descriptor instance.

**static bind\_early** (*instance*)

Bind all *BoundClass* attributes of the *instance*'s class to the instance itself to increase performance.

**class** `simpy.core.EmptySchedule`

Thrown by an *Environment* if there are no further events to be processed.

`simpy.core.Infinity = inf`

Convenience alias for infinity

## simpy.exceptions — Exception types used by SimPy

SimPy specific exceptions.

**exception** `simpy.exceptions.SimPyException`

Base class for all SimPy specific exceptions.

**exception** `simpy.exceptions.Interrupt` (*cause*)

Exception thrown into a process if it is interrupted (see *interrupt()*).

*cause* provides the reason for the interrupt, if any.

If a process is interrupted concurrently, all interrupts will be thrown into the process in the same order as they occurred.

**cause**

The cause of the interrupt or None if no cause was provided.

**exception** `simpy.exceptions.StopProcess` (*value*)

Raised to stop a SimPy process (similar to *StopIteration*).

In Python 2, a return value inside generator functions is not allowed. The fall-back was raising *StopIteration*(*value*) instead. However, this is deprecated now, so we need a custom exception type for this.

**value**

The process' return value.

## simpy.events — Core event types

This module contains the basic event types used in SimPy.

The base class for all events is *Event*. Though it can be directly used, there are several specialized subclasses of it.

<i>Event</i> ( <i>env</i> )	An event that may happen at some point in time.
<i>Timeout</i> ( <i>env</i> , <i>delay</i> [, <i>value</i> ])	A <i>Event</i> that gets triggered after a <i>delay</i> has passed.
<i>Process</i> ( <i>env</i> , <i>generator</i> )	Process an event yielding generator.
<i>AnyOf</i> ( <i>env</i> , <i>events</i> )	A <i>Condition</i> event that is triggered if any of a list of <i>events</i> has been successfully triggered.
<i>AllOf</i> ( <i>env</i> , <i>events</i> )	A <i>Condition</i> event that is triggered if all of a list of <i>events</i> have been successfully triggered.

`simpy.events.PENDING = object()`

Unique object to identify pending values of events.

`simpy.events.URGENT = 0`

Priority of interrupts and process initialization events.

`simpy.events.NORMAL = 1`

Default priority used by events.

**class** `simpy.events.Event` (*env*)

An event that may happen at some point in time.

An event

- may happen (*triggered* is `False`),
- is going to happen (*triggered* is `True`) or
- has happened (*processed* is `True`).

Every event is bound to an environment *env* and is initially not triggered. Events are scheduled for processing by the environment after they are triggered by either `succeed()`, `fail()` or `trigger()`. These methods also set the *ok* flag and the *value* of the event.

An event has a list of *callbacks*. A callback can be any callable. Once an event gets processed, all callbacks will be invoked with the event as the single argument. Callbacks can check if the event was successful by examining *ok* and do further processing with the *value* it has produced.

Failed events are never silently ignored and will raise an exception upon being processed. If a callback handles an exception, it must set *defused* to `True` to prevent this.

This class also implements `__and__()` (`&`) and `__or__()` (`|`). If you concatenate two events using one of these operators, a *Condition* event is generated that lets you wait for both or one of them.

**env = None**

The *Environment* the event lives in.

**callbacks = None**

List of functions that are called when the event is processed.

**triggered**

Becomes `True` if the event has been triggered and its callbacks are about to be invoked.

**processed**

Becomes `True` if the event has been processed (e.g., its callbacks have been invoked).

**ok**

Becomes `True` when the event has been triggered successfully.

A “successful” event is one triggered with `succeed()`.

**Levanta `AttributeError`** – if accessed before the event is triggered.

**defused**

Becomes `True` when the failed event’s exception is “defused”.

When an event fails (i.e. with `fail()`), the failed event’s *value* is an exception that will be re-raised when the *Environment* processes the event (i.e. in `step()`).

It is also possible for the failed event’s exception to be defused by setting *defused* to `True` from an event callback. Doing so prevents the event’s exception from being re-raised when the event is processed by the *Environment*.

**value**

The value of the event if it is available.

The value is available when the event has been triggered.

Raises `AttributeError` if the value is not yet available.

**trigger** (*event*)

Trigger the event with the state and value of the provided *event*. Return *self* (this event instance).

This method can be used directly as a callback function to trigger chain reactions.

**succeed** (*value=None*)

Set the event's value, mark it as successful and schedule it for processing by the environment. Returns the event instance.

Raises `RuntimeError` if this event has already been triggered.

**fail** (*exception*)

Set *exception* as the events value, mark it as failed and schedule it for processing by the environment. Returns the event instance.

Raises `ValueError` if *exception* is not an `Exception`.

Raises `RuntimeError` if this event has already been triggered.

**class** `simpy.events.Timeout` (*env, delay, value=None*)

A `Event` that gets triggered after a *delay* has passed.

This event is automatically triggered when it is created.

**defused**

Becomes `True` when the failed event's exception is "defused".

When an event fails (i.e. with `fail()`), the failed event's *value* is an exception that will be re-raised when the `Environment` processes the event (i.e. in `step()`).

It is also possible for the failed event's exception to be defused by setting `defused` to `True` from an event callback. Doing so prevents the event's exception from being re-raised when the event is processed by the `Environment`.

**fail** (*exception*)

Set *exception* as the events value, mark it as failed and schedule it for processing by the environment. Returns the event instance.

Raises `ValueError` if *exception* is not an `Exception`.

Raises `RuntimeError` if this event has already been triggered.

**ok**

Becomes `True` when the event has been triggered successfully.

A "successful" event is one triggered with `succeed()`.

**Levanta `AttributeError`** – if accessed before the event is triggered.

**processed**

Becomes `True` if the event has been processed (e.g., its callbacks have been invoked).

**succeed** (*value=None*)

Set the event's value, mark it as successful and schedule it for processing by the environment. Returns the event instance.

Raises `RuntimeError` if this event has already been triggered.

**trigger** (*event*)

Trigger the event with the state and value of the provided *event*. Return `self` (this event instance).

This method can be used directly as a callback function to trigger chain reactions.

**triggered**

Becomes `True` if the event has been triggered and its callbacks are about to be invoked.

**value**

The value of the event if it is available.

The value is available when the event has been triggered.

Raises `AttributeError` if the value is not yet available.

**class** `simpy.events.Initialize` (*env, process*)  
 Initializes a process. Only used internally by `Process`.

This event is automatically triggered when it is created.

**defused**

Becomes `True` when the failed event's exception is "defused".

When an event fails (i.e. with `fail()`), the failed event's *value* is an exception that will be re-raised when the `Environment` processes the event (i.e. in `step()`).

It is also possible for the failed event's exception to be defused by setting `defused` to `True` from an event callback. Doing so prevents the event's exception from being re-raised when the event is processed by the `Environment`.

**fail** (*exception*)

Set *exception* as the events value, mark it as failed and schedule it for processing by the environment. Returns the event instance.

Raises `ValueError` if *exception* is not an `Exception`.

Raises `RuntimeError` if this event has already been triggered.

**ok**

Becomes `True` when the event has been triggered successfully.

A "successful" event is one triggered with `succeed()`.

**Levanta `AttributeError`** – if accessed before the event is triggered.

**processed**

Becomes `True` if the event has been processed (e.g., its callbacks have been invoked).

**succeed** (*value=None*)

Set the event's value, mark it as successful and schedule it for processing by the environment. Returns the event instance.

Raises `RuntimeError` if this event has already been triggered.

**trigger** (*event*)

Trigger the event with the state and value of the provided *event*. Return *self* (this event instance).

This method can be used directly as a callback function to trigger chain reactions.

**triggered**

Becomes `True` if the event has been triggered and its callbacks are about to be invoked.

**value**

The value of the event if it is available.

The value is available when the event has been triggered.

Raises `AttributeError` if the value is not yet available.

**class** `simpy.events.Interruption` (*process, cause*)  
 Immediately schedules an `Interrupt` exception with the given *cause* to be thrown into *process*.

This event is automatically triggered when it is created.

**defused**

Becomes `True` when the failed event's exception is "defused".

When an event fails (i.e. with `fail()`), the failed event's *value* is an exception that will be re-raised when the *Environment* processes the event (i.e. in `step()`).

It is also possible for the failed event's exception to be defused by setting `defused` to `True` from an event callback. Doing so prevents the event's exception from being re-raised when the event is processed by the *Environment*.

**fail** (*exception*)

Set *exception* as the events value, mark it as failed and schedule it for processing by the environment. Returns the event instance.

Raises `ValueError` if *exception* is not an `Exception`.

Raises `RuntimeError` if this event has already been triggered.

**ok**

Becomes `True` when the event has been triggered successfully.

A “successful” event is one triggered with `succeed()`.

Levanta `AttributeError` – if accessed before the event is triggered.

**processed**

Becomes `True` if the event has been processed (e.g., its callbacks have been invoked).

**succeed** (*value=None*)

Set the event's value, mark it as successful and schedule it for processing by the environment. Returns the event instance.

Raises `RuntimeError` if this event has already been triggered.

**trigger** (*event*)

Trigger the event with the state and value of the provided *event*. Return *self* (this event instance).

This method can be used directly as a callback function to trigger chain reactions.

**triggered**

Becomes `True` if the event has been triggered and its callbacks are about to be invoked.

**value**

The value of the event if it is available.

The value is available when the event has been triggered.

Raises `AttributeError` if the value is not yet available.

**class** `simpy.events.Process` (*env, generator*)

Process an event yielding generator.

A generator (also known as a coroutine) can suspend its execution by yielding an event. `Process` will take care of resuming the generator with the value of that event once it has happened. The exception of failed events is thrown into the generator.

`Process` itself is an event, too. It is triggered, once the generator returns or raises an exception. The value of the process is the return value of the generator or the exception, respectively.

---

**Nota:** Python version prior to 3.3 do not support return statements in generators. You can use `:meth:~simpy.core.Environment.exit()` as a workaround.

---

Processes can be interrupted during their execution by `interrupt()`.

**target**

The event that the process is currently waiting for.

Returns `None` if the process is dead or it is currently being interrupted.

**is\_alive**

True until the process generator exits.

**interrupt** (*cause=None*)

Interrupt this process optionally providing a *cause*.

A process cannot be interrupted if it already terminated. A process can also not interrupt itself. Raise a `RuntimeError` in these cases.

**defused**

Becomes `True` when the failed event's exception is "defused".

When an event fails (i.e. with `fail()`), the failed event's *value* is an exception that will be re-raised when the `Environment` processes the event (i.e. in `step()`).

It is also possible for the failed event's exception to be defused by setting `defused` to `True` from an event callback. Doing so prevents the event's exception from being re-raised when the event is processed by the `Environment`.

**fail** (*exception*)

Set *exception* as the events value, mark it as failed and schedule it for processing by the environment. Returns the event instance.

Raises `ValueError` if *exception* is not an `Exception`.

Raises `RuntimeError` if this event has already been triggered.

**ok**

Becomes `True` when the event has been triggered successfully.

A "successful" event is one triggered with `succeed()`.

**Levanta `AttributeError`** – if accessed before the event is triggered.

**processed**

Becomes `True` if the event has been processed (e.g., its callbacks have been invoked).

**succeed** (*value=None*)

Set the event's value, mark it as successful and schedule it for processing by the environment. Returns the event instance.

Raises `RuntimeError` if this event has already been triggered.

**trigger** (*event*)

Trigger the event with the state and value of the provided *event*. Return *self* (this event instance).

This method can be used directly as a callback function to trigger chain reactions.

**triggered**

Becomes `True` if the event has been triggered and its callbacks are about to be invoked.

**value**

The value of the event if it is available.

The value is available when the event has been triggered.

Raises `AttributeError` if the value is not yet available.

**class** `simpy.events.Condition` (*env, evaluate, events*)

An event that gets triggered once the condition function *evaluate* returns `True` on the given list of *events*.

The value of the condition event is an instance of `ConditionValue` which allows convenient access to the input events and their values. The `ConditionValue` will only contain entries for those events that occurred before the condition is processed.

If one of the events fails, the condition also fails and forwards the exception of the failing event.

The `evaluate` function receives the list of target events and the number of processed events in this list: `evaluate(events, processed_count)`. If it returns `True`, the condition is triggered. The `Condition.all_events()` and `Condition.any_events()` functions are used to implement *and* (&) and *or* (|) for events.

Condition events can be nested.

**static all\_events** (*events, count*)

An evaluation function that returns `True` if all *events* have been triggered.

**static any\_events** (*events, count*)

An evaluation function that returns `True` if at least one of *events* has been triggered.

**defused**

Becomes `True` when the failed event's exception is "defused".

When an event fails (i.e. with `fail()`), the failed event's *value* is an exception that will be re-raised when the `Environment` processes the event (i.e. in `step()`).

It is also possible for the failed event's exception to be defused by setting `defused` to `True` from an event callback. Doing so prevents the event's exception from being re-raised when the event is processed by the `Environment`.

**fail** (*exception*)

Set *exception* as the events value, mark it as failed and schedule it for processing by the environment. Returns the event instance.

Raises `ValueError` if *exception* is not an `Exception`.

Raises `RuntimeError` if this event has already been triggered.

**ok**

Becomes `True` when the event has been triggered successfully.

A "successful" event is one triggered with `succeed()`.

**Levanta `AttributeError`** – if accessed before the event is triggered.

**processed**

Becomes `True` if the event has been processed (e.g., its callbacks have been invoked).

**succeed** (*value=None*)

Set the event's value, mark it as successful and schedule it for processing by the environment. Returns the event instance.

Raises `RuntimeError` if this event has already been triggered.

**trigger** (*event*)

Trigger the event with the state and value of the provided *event*. Return *self* (this event instance).

This method can be used directly as a callback function to trigger chain reactions.

**triggered**

Becomes `True` if the event has been triggered and its callbacks are about to be invoked.

**value**

The value of the event if it is available.

The value is available when the event has been triggered.

Raises `AttributeError` if the value is not yet available.

**class** `simpy.events.AllOf` (*env, events*)

A *Condition* event that is triggered if all of a list of *events* have been successfully triggered. Fails immediately if any of *events* failed.

**defused**

Becomes `True` when the failed event's exception is "defused".

When an event fails (i.e. with `fail()`), the failed event's *value* is an exception that will be re-raised when the *Environment* processes the event (i.e. in `step()`).

It is also possible for the failed event's exception to be defused by setting `defused` to `True` from an event callback. Doing so prevents the event's exception from being re-raised when the event is processed by the *Environment*.

**fail** (*exception*)

Set *exception* as the events value, mark it as failed and schedule it for processing by the environment. Returns the event instance.

Raises `ValueError` if *exception* is not an `Exception`.

Raises `RuntimeError` if this event has already been triggered.

**ok**

Becomes `True` when the event has been triggered successfully.

A "successful" event is one triggered with `succeed()`.

**Levanta `AttributeError`** – if accessed before the event is triggered.

**processed**

Becomes `True` if the event has been processed (e.g., its callbacks have been invoked).

**succeed** (*value=None*)

Set the event's value, mark it as successful and schedule it for processing by the environment. Returns the event instance.

Raises `RuntimeError` if this event has already been triggered.

**trigger** (*event*)

Trigger the event with the state and value of the provided *event*. Return *self* (this event instance).

This method can be used directly as a callback function to trigger chain reactions.

**triggered**

Becomes `True` if the event has been triggered and its callbacks are about to be invoked.

**value**

The value of the event if it is available.

The value is available when the event has been triggered.

Raises `AttributeError` if the value is not yet available.

**class** `simpy.events.AnyOf` (*env, events*)

A *Condition* event that is triggered if any of a list of *events* has been successfully triggered. Fails immediately if any of *events* failed.

**defused**

Becomes `True` when the failed event's exception is "defused".

When an event fails (i.e. with `fail()`), the failed event's *value* is an exception that will be re-raised when the *Environment* processes the event (i.e. in `step()`).

It is also possible for the failed event's exception to be defused by setting `defused` to `True` from an event callback. Doing so prevents the event's exception from being re-raised when the event is processed by the `Environment`.

**fail** (*exception*)

Set *exception* as the events value, mark it as failed and schedule it for processing by the environment. Returns the event instance.

Raises `ValueError` if *exception* is not an `Exception`.

Raises `RuntimeError` if this event has already been triggered.

**ok**

Becomes `True` when the event has been triggered successfully.

A "successful" event is one triggered with `succeed()`.

Levanta `AttributeError` – if accessed before the event is triggered.

**processed**

Becomes `True` if the event has been processed (e.g., its callbacks have been invoked).

**succeed** (*value=None*)

Set the event's value, mark it as successful and schedule it for processing by the environment. Returns the event instance.

Raises `RuntimeError` if this event has already been triggered.

**trigger** (*event*)

Trigger the event with the state and value of the provided *event*. Return *self* (this event instance).

This method can be used directly as a callback function to trigger chain reactions.

**triggered**

Becomes `True` if the event has been triggered and its callbacks are about to be invoked.

**value**

The value of the event if it is available.

The value is available when the event has been triggered.

Raises `AttributeError` if the value is not yet available.

**class** `simpy.events.ConditionValue`

Result of a `Condition`. It supports convenient dict-like access to the triggered events and their values. The events are ordered by their occurrences in the condition.

## simpy.resources — Shared resource primitives

SimPy implements three types of resources that can be used to synchronize processes or to model congestion points:

<code>resource</code>	Shared resources supporting priorities and preemption.
<code>container</code>	Resource for sharing homogeneous matter between processes, either continuous (like water) or discrete (like apples).
<code>store</code>	Shared resources for storing a possibly unlimited amount of objects supporting requests for specific objects.

They are derived from the base classes defined in the `base` module. These classes are also meant to support the implementation of custom resource types.

## Resources — `simpy.resources.resource`

Shared resources supporting priorities and preemption.

These resources can be used to limit the number of processes using them concurrently. A process needs to *request* the usage right to a resource. Once the usage right is not needed anymore it has to be *released*. A gas station can be modelled as a resource with a limited amount of fuel-pumps. Vehicles arrive at the gas station and request to use a fuel-pump. If all fuel-pumps are in use, the vehicle needs to wait until one of the users has finished refueling and releases its fuel-pump.

These resources can be used by a limited number of processes at a time. Processes *request* these resources to become a user and have to *release* them once they are done. For example, a gas station with a limited number of fuel pumps can be modeled with a *Resource*. Arriving vehicles request a fuel-pump. Once one is available they refuel. When they are done, they release the fuel-pump and leave the gas station.

Requesting a resource is modelled as “putting a process’ token into the resources” and releasing a resource correspondingly as “getting a process’ token out of the resource”. Thus, calling `request()/release()` is equivalent to calling `put()/get()`. Note, that releasing a resource will always succeed immediately, no matter if a process is actually using a resource or not.

Besides *Resource*, there is a *PriorityResource*, where processes can define a request priority, and a *PreemptiveResource* whose resource users can be preempted by requests with a higher priority.

**class** `simpy.resources.resource.Resource` (*env*, *capacity=1*)

Resource with *capacity* of usage slots that can be requested by processes.

If all slots are taken, requests are enqueued. Once a usage request is released, a pending request will be triggered.

The *env* parameter is the *Environment* instance the resource is bound to.

**users = None**

List of *Request* events for the processes that are currently using the resource.

**queue = None**

Queue of pending *Request* events. Alias of *put\_queue*.

**count**

Number of users currently using the resource.

**request**

Request a usage slot.

apelido de *Request*

**release**

Release a usage slot.

apelido de *Release*

**class** `simpy.resources.resource.PriorityResource` (*env*, *capacity=1*)

A *Resource* supporting prioritized requests.

Pending requests in the *queue* are sorted in ascending order by their *priority* (that means lower values are more important).

**PutQueue**

Type of the put queue. See *put\_queue* for details.

apelido de *SortedQueue*

**GetQueue**

Type of the get queue. See *get\_queue* for details.

apelido de *list*

**request**

Request a usage slot with the given *priority*.

apelido de *PriorityRequest*

**release**

Release a usage slot.

apelido de *Release*

**class** `simpy.resources.resource.PreemptiveResource` (*env, capacity=1*)  
 A *PriorityResource* with preemption.

If a request is preempted, the process of that request will receive an *Interrupt* with a *Preempted* instance as cause.

**class** `simpy.resources.resource.Preempted` (*by, usage\_since, resource*)  
 Cause of an preemption *Interrupt* containing information about the preemption.

**by = None**

The preempting *simpy.events.Process*.

**usage\_since = None**

The simulation time at which the preempted process started to use the resource.

**resource = None**

The resource which was lost, i.e., caused the preemption.

**class** `simpy.resources.resource.Request` (*resource*)  
 Request usage of the *resource*. The event is triggered once access is granted. Subclass of *simpy.resources.base.Put*.

If the maximum capacity of users has not yet been reached, the request is triggered immediately. If the maximum capacity has been reached, the request is triggered once an earlier usage request on the resource is released.

The request is automatically released when the request was created within a `with` statement.

**class** `simpy.resources.resource.PriorityRequest` (*resource, priority=0, preempt=True*)  
 Request the usage of *resource* with a given *priority*. If the *resource* supports preemption and *preempt* is `True` other usage requests of the *resource* may be preempted (see *PreemptiveResource* for details).

This event type inherits *Request* and adds some additional attributes needed by *PriorityResource* and *PreemptiveResource*

**priority = None**

The priority of this request. A smaller number means higher priority.

**preempt = None**

Indicates whether the request should preempt a resource user or not (*PriorityResource* ignores this flag).

**time = None**

The time at which the request was made.

**usage\_since = None**

The time at which the request succeeded.

**key = None**

Key for sorting events. Consists of the priority (lower value is more important), the time at which the request was made (earlier requests are more important) and finally the preemption flag (preempt requests are more important).

**class** `simpy.resources.resource.Release(resource, request)`  
 Releases the usage of *resource* granted by *request*. This event is triggered immediately. Subclass of `simpy.resources.base.Get`.

**request = None**

The request (*Request*) that is to be released.

**class** `simpy.resources.resource.SortedQueue(maxlen=None)`  
 Queue for sorting events by their *key* attribute.

**maxlen = None**

Maximum length of the queue.

**append(item)**

Sort *item* into the queue.

Raise a `RuntimeError` if the queue is full.

## Containers — `simpy.resources.container`

Resource for sharing homogeneous matter between processes, either continuous (like water) or discrete (like apples).

A *Container* can be used to model the fuel tank of a gasoline station. Tankers increase and refuelled cars decrease the amount of gas in the station's fuel tanks.

**class** `simpy.resources.container.Container(env, capacity=inf, init=0)`

Resource containing up to *capacity* of matter which may either be continuous (like water) or discrete (like apples). It supports requests to put or get matter into/from the container.

The *env* parameter is the *Environment* instance the container is bound to.

The *capacity* defines the size of the container. By default, a container is of unlimited size. The initial amount of matter is specified by *init* and defaults to 0.

Raise a `ValueError` if `capacity <= 0, init < 0 or init > capacity`.

**level**

The current amount of the matter in the container.

**put**

Request to put *amount* of matter into the container.

apelido de *ContainerPut*

**get**

Request to get *amount* of matter out of the container.

apelido de *ContainerGet*

**class** `simpy.resources.container.ContainerPut(container, amount)`

Request to put *amount* of matter into the *container*. The request will be triggered once there is enough space in the *container* available.

Raise a `ValueError` if `amount <= 0`.

**amount = None**

The amount of matter to be put into the container.

**class** `simpy.resources.container.ContainerGet(container, amount)`

Request to get *amount* of matter from the *container*. The request will be triggered once there is enough matter available in the *container*.

Raise a `ValueError` if `amount <= 0`.

**amount = None**

The amount of matter to be taken out of the container.

## Stores — `simpy.resources.store`

Shared resources for storing a possibly unlimited amount of objects supporting requests for specific objects.

The `Store` operates in a FIFO (first-in, first-out) order. Objects are retrieved from the store in the order they were put in. The `get` requests of a `FilterStore` can be customized by a filter to only retrieve objects matching a given criterion.

**class** `simpy.resources.store.Store` (*env*, *capacity=inf*)

Resource with *capacity* slots for storing arbitrary objects. By default, the *capacity* is unlimited and objects are put and retrieved from the store in a first-in first-out order.

The *env* parameter is the `Environment` instance the container is bound to.

**items = None**

List of the items available in the store.

**put**

Request to put *item* into the store.

apelido de `StorePut`

**get**

Request to get an *item* out of the store.

apelido de `StoreGet`

**class** `simpy.resources.store.PriorityItem` (*priority*, *item*)

Wrap an arbitrary *item* with an orderable *priority*.

Pairs a *priority* with an arbitrary *item*. Comparisons of `PriorityItem` instances only consider the *priority* attribute, thus supporting use of unorderable items in a `PriorityStore` instance.

**item**

Alias for field number 1

**priority**

Alias for field number 0

**class** `simpy.resources.store.PriorityStore` (*env*, *capacity=inf*)

Resource with *capacity* slots for storing objects in priority order.

Unlike `Store` which provides first-in first-out discipline, `PriorityStore` maintains items in sorted order such that the smallest items value are retrieved first from the store.

All items in a `PriorityStore` instance must be orderable; which is to say that items must implement `__lt__()`. To use unorderable items with `PriorityStore`, use `PriorityItem`.

**class** `simpy.resources.store.FilterStore` (*env*, *capacity=inf*)

Resource with *capacity* slots for storing arbitrary objects supporting filtered get requests. Like the `Store`, the *capacity* is unlimited by default and objects are put and retrieved from the store in a first-in first-out order.

Get requests can be customized with a filter function to only trigger for items for which said filter function returns `True`.

---

**Nota:** In contrast to `Store`, get requests of a `FilterStore` won't necessarily be triggered in the same order they were issued.

*Example:* The store is empty. *Process 1* tries to get an item of type *a*, *Process 2* an item of type *b*. Another process puts one item of type *b* into the store. Though *Process 2* made his request after *Process 1*, it will receive that new item because *Process 1* doesn't want it.

**put**

Request a to put *item* into the store.

apelido de *StorePut*

**get**

Request a to get an *item*, for which *filter* returns `True`, out of the store.

apelido de *FilterStoreGet*

**class** `simpy.resources.store.StorePut` (*store, item*)

Request to put *item* into the *store*. The request is triggered once there is space for the item in the store.

**item = None**

The item to put into the store.

**class** `simpy.resources.store.StoreGet` (*resource*)

Request to get an *item* from the *store*. The request is triggered once there is an item available in the store.

**class** `simpy.resources.store.FilterStoreGet` (*resource, filter=<function <lambda>>*)

Request to get an *item* from the *store* matching the *filter*. The request is triggered once there is such an item available in the store.

*filter* is a function receiving one item. It should return `True` for items matching the filter criterion. The default function returns `True` for all items, which makes the request to behave exactly like *StoreGet*.

**filter = None**

The filter function to filter items in the store.

## Base classes — `simpy.resources.base`

Base classes of for SimPy's shared resource types.

*BaseResource* defines the abstract base resource. It supports *get* and *put* requests, which return *Put* and *Get* events respectively. These events are triggered once the request has been completed.

**class** `simpy.resources.base.BaseResource` (*env, capacity*)

Abstract base class for a shared resource.

You can *put* () something into the resources or *get* () something out of it. Both methods return an event that is triggered once the operation is completed. If a *put* () request cannot complete immediately (for example if the resource has reached a capacity limit) it is enqueued in the *put\_queue* for later processing. Likewise for *get* () requests.

Subclasses can customize the resource by:

- providing custom *PutQueue* and *GetQueue* types,
- providing custom *Put* respectively *Get* events,
- and implementing the request processing behaviour through the methods `_do_get` () and `_do_put` () .

**PutQueue**

The type to be used for the *put\_queue*. It is a plain *list* by default. The type must support index access (e.g. `__getitem__` () and `__len__` ()) as well as provide `append` () and `pop` () operations.

apelido de *list*

**GetQueue**

The type to be used for the `get_queue`. It is a plain `list` by default. The type must support index access (e.g. `__getitem__()` and `__len__()`) as well as provide `append()` and `pop()` operations.

apelido de `list`

**put\_queue = None**

Queue of pending `put` requests.

**get\_queue = None**

Queue of pending `get` requests.

**capacity**

Maximum capacity of the resource.

**put**

Request to put something into the resource and return a `Put` event, which gets triggered once the request succeeds.

apelido de `Put`

**get**

Request to get something from the resource and return a `Get` event, which gets triggered once the request succeeds.

apelido de `Get`

**class** `simpy.resources.base.Put(resource)`

Generic event for requesting to put something into the `resource`.

This event (and all of its subclasses) can act as context manager and can be used with the `with` statement to automatically cancel the request if an exception (like an `simpy.exceptions.Interrupt` for example) occurs:

```
with res.put(item) as request:
    yield request
```

**cancel()**

Cancel this put request.

This method has to be called if the put request must be aborted, for example if a process needs to handle an exception like an `Interrupt`.

If the put request was created in a `with` statement, this method is called automatically.

**class** `simpy.resources.base.Get(resource)`

Generic event for requesting to get something from the `resource`.

This event (and all of its subclasses) can act as context manager and can be used with the `with` statement to automatically cancel the request if an exception (like an `simpy.exceptions.Interrupt` for example) occurs:

```
with res.get() as request:
    item = yield request
```

**cancel()**

Cancel this get request.

This method has to be called if the get request must be aborted, for example if a process needs to handle an exception like an `Interrupt`.

If the get request was created in a `with` statement, this method is called automatically.

## simpy.rt — Real-time simulation

Execution environment for events that synchronizes passing of time with the real-time (aka *wall-clock time*).

**class** `simpy.rt.RealtimeEnvironment` (*initial\_time=0, factor=1.0, strict=True*)

Execution environment for an event-based simulation which is synchronized with the real-time (also known as wall-clock time). A time step will take *factor* seconds of real time (one second by default). A step from 0 to 3 with a *factor=0.5* will, for example, take at least 1.5 seconds.

The `step()` method will raise a `RuntimeError` if a time step took too long to compute. This behaviour can be disabled by setting *strict* to `False`.

**now**

The current simulation time.

**active\_process**

The currently active process of the environment.

**factor**

Scaling factor of the real-time.

**strict**

Running mode of the environment. `step()` will raise a `RuntimeError` if this is set to `True` and the processing of events takes too long.

**process** (*generator*)

Create a new `Process` instance for *generator*.

**timeout** (*delay, value=None*)

Return a new `Timeout` event with a *delay* and, optionally, a *value*.

**event** ()

Return a new `Event` instance. Yielding this event suspends a process until another process triggers the event.

**all\_of** (*events*)

Return a new `AllOf` condition for a list of *events*.

**any\_of** (*events*)

Return a new `AnyOf` condition for a list of *events*.

**exit** (*value=None*)

Stop the current process, optionally providing a *value*.

This is a convenience function provided for Python versions prior to 3.3. From Python 3.3, you can instead use `return value` in a process.

**schedule** (*event, priority=1, delay=0*)

Schedule an *event* with a given *priority* and a *delay*.

**peek** ()

Get the time of the next scheduled event. Return `Infinity` if there is no further event.

**step** ()

Process the next event after enough real-time has passed for the event to happen.

The delay is scaled according to the real-time *factor*. With *strict* mode enabled, a `RuntimeError` will be raised, if the event is processed too slowly.

**sync** ()

Synchronize the internal time with the current wall-clock time.

This can be useful to prevent `step()` from raising an error if a lot of time passes between creating the `RealtimeEnvironment` and calling `run()` or `step()`.

**run** (*until=None*)

Executes `step()` until the given criterion *until* is met.

- If it is `None` (which is the default), this method will return when there are no further events to be processed.
- If it is an `Event`, the method will continue stepping until this event has been triggered and will return its value. Raises a `RuntimeError` if there are no further events to be processed and the *until* event was not triggered.
- If it is a number, the method will continue stepping until the environment's time reaches *until*.

## simpy.util — Utility functions for SimPy

A collection of utility functions:

---

`start_delayed(env, generator, delay)` Return a helper process that starts another process for *generator* after a certain *delay*.

---

`simpy.util.start_delayed(env, generator, delay)`

Return a helper process that starts another process for *generator* after a certain *delay*.

`process()` starts a process at the current simulation time. This helper allows you to start a process after a delay of *delay* simulation time units:

```
>>> from simpy import Environment
>>> from simpy.util import start_delayed
>>> def my_process(env, x):
...     print('%s, %s' % (env.now, x))
...     yield env.timeout(1)
...
>>> env = Environment()
>>> proc = start_delayed(env, my_process(env, 3), 5)
>>> env.run()
5, 3
```

Raise a `ValueError` if `delay <= 0`.

---

## About SimPy

---

This sections is all about the non-technical stuff. How did SimPy evolve? Who was responsible for it? And what the heck were they thinking when they made it?

### SimPy History & Change Log

SimPy was originally based on ideas from Simula and Simscript but uses standard Python. It combines two previous packages, SiPy, in Simula-Style (Klaus Müller) and SimPy, in Simscript style (Tony Vignaux and Chang Chui).

SimPy was based on efficient implementation of co-routines using Python's generators capability.

SimPy 3 introduced a completely new and easier-to-use API, but still relied on Python's generators as they proved to work very well.

The package has been hosted on Sourceforge.net since September 15th, 2002. In June 2012, the project moved to Bitbucket.org.

#### 3.0.9 – 2016-06-12

- [NEW] *PriorityStore* resource and performance benchmarks were implemented by Peter Grayson.
- [FIX] Support for identifying nested preemptions was added by Cristian Klein.

#### 3.0.8 – 2015-06-23

- [NEW] Added a monitoring guide to the documentation.
- [FIX] Improved packaging (thanks to Larissa Reis).
- [FIX] Fixed and improved various test cases.

#### 3.0.7 - 2015-03-01

- [FIX] State of resources and requests were inconsistent before the request has been processed (*issue #62*).
- [FIX] Empty conditions were never triggered (regression in 3.0.6, *issue #63*).
- [FIX] `Environment.run()` will fail if the until event does not get triggered (*issue #64*).
- [FIX] Callback modification during event processing is now prohibited (thanks to Andreas Beham).

### 3.0.6 - 2015-01-30

- [NEW] Guide to SimPy resources.
- [CHANGE] Improve performance of condition events.
- [CHANGE] Improve performance of filter store (thanks to Christoph Körner).
- [CHANGE] Exception tracebacks are now more compact.
- [FIX] `Allof` conditions handle already processed events correctly (issue #52).
- [FIX] Add `sync()` to `RealtimeEnvironment` to reset its internal wall-clock reference time (issue #42).
- [FIX] Only send copies of exceptions into processes to prevent traceback modifications.
- [FIX] Documentation improvements.

### 3.0.5 – 2014-05-14

- [CHANGE] Move interruption and all of the safety checks into a new event (pull request #30)
- [FIX] `FilterStore.get()` now behaves correctly (issue #49).
- [FIX] Documentation improvements.

### 3.0.4 – 2014-04-07

- [NEW] Verified, that SimPy works on Python 3.4.
- [NEW] Guide to SimPy events
- [CHANGE] The result dictionary for condition events (`Allof / &` and `AnyOf / |`) now is an *OrderedDict* sorted in the same way as the original events list.
- [CHANGE] Condition events now also except processed events.
- [FIX] `Resource.request()` directly after `Resource.release()` no longer successful. The process now has to wait as supposed to.
- [FIX] `Event.fail()` now accept all exceptions derived from `BaseException` instead of only `Exception`.

### 3.0.3 – 2014-03-06

- [NEW] Guide to SimPy basics.
- [NEW] Guide to SimPy Environments.
- [FIX] Timing problems with real time simulation on Windows (issue #46).
- [FIX] Installation problems on Windows due to Unicode errors (issue #41).
- [FIX] Minor documentation issues.

### 3.0.2 – 2013-10-24

- [FIX] The default capacity for `Container` and `FilterStore` is now also `inf`.

### 3.0.1 – 2013-10-24

- [FIX] Documentation and default parameters of `Store` didn't match. Its default capacity is now `inf`.

## 3.0 – 2013-10-11

SimPy 3 has been completely rewritten from scratch. Our main goals were to simplify the API and code base as well as making SimPy more flexible and extensible. Some of the most important changes are:

- Stronger focus on events. Processes yield event instances and are suspended until the event is triggered. An example for an event is a *timeout* (formerly known as *hold*), but even processes are now events, too (you can wait until a process terminates).
- Events can be combined with `&` (and) and `|` (or) to create *condition events*.
- Process can now be defined by any generator function. You don't have to subclass `Process` anymore.
- No more global simulation state. Every simulation stores its state in an *environment* which is comparable to the old `Simulation` class.
- Improved resource system with newly added resource types.
- Removed plotting and GUI capabilities. `Pyside` and `matplotlib` are much better with this.
- Greatly improved test suite. Its cleaner, and the tests are shorter and more numerous.
- Completely overhauled documentation.

There is a [guide for porting from SimPy 2 to SimPy 3](#). If you want to stick to SimPy 2 for a while, change your requirements to `'SimPy>=2.3,<3'`.

All in all, SimPy has become a framework for asynchronous programming based on coroutines. It brings more than ten years of experience and scientific know-how in the field of event-discrete simulation to the world of asynchronous programming and should thus be a solid foundation for everything based on an event loop.

You can find information about older versions on the [history page](#)

### 2.3.1 – 2012-01-28

- [NEW] More improvements on the documentation.
- [FIX] Syntax error in `tkconsole.py` when installing on Py3.2.
- [FIX] Added *mock* to the dep. list in `SimPy.test()`.

## 2.3 – 2011-12-24

- [NEW] Support for Python 3.2. Support for Python `<= 2.5` has been dropped.
- [NEW] `SimPy.test()` method to run the tests on the installed version of SimPy.
- [NEW] Tutorials/examples were integrated into the test suite.
- [CHANGE] Even more code clean-up (e.g., removed prints throughout the code, removed if-main-blocks, ...).
- [CHANGE] Many documentation improvements.

## 2.2 – 2011-09-27

- [CHANGE] Restructured package layout to be conform to the [Hitchhiker's Guide to packaging](#)
- [CHANGE] Tests have been ported to pytest.
- [CHANGE] Documentation improvements and clean-ups.
- [FIX] Fixed incorrect behavior of `Store._put`, thanks to Johannes Koomer for the fix.

## 2.1 – 2010-06-03

- [NEW] A function *step* has been added to the API. When called, it executes the next scheduled event. (*step* is actually a method of *Simulation*.)
- [NEW] Another new function is *peek*. It returns the time of the next event. By using *peek* and *step* together, one can easily write e.g. an interactive program to step through a simulation event by event.
- [NEW] A simple interactive debugger `stepping.py` has been added. It allows stepping through a simulation, with options to skip to a certain time, skip to the next event of a given process, or viewing the event list.
- [NEW] Versions of the Bank tutorials (documents and programs) using the advanced- [NEW] object-oriented API have been added.
- [NEW] A new document describes tools for gaining insight into and debugging SimPy models.
- [CHANGE] Major re-structuring of SimPy code, resulting in much less SimPy code – great for the maintainers.
- [CHANGE] Checks have been added which test whether entities belong to the same Simulation instance.
- [CHANGE] The Monitor and Tally methods `timeAverage` and `timeVariance` now calculate only with the observed time-series. No value is assumed for the period prior to the first observation.
- [CHANGE] Changed class `Lister` so that circular references between objects no longer lead to stack overflow and crash.
- [FIX] Functions `allEventNotices` and `allEventTimes` are working again.
- [FIX] Error messages for methods in `SimPy.Lib` work again.

### 2.0.1 – 2009-04-06

- [NEW] Tests for real time behavior (`testRT_Behavior.py` and `testRT_Behavior_OO.py` in folder `SimPy`).
- [FIX] Repaired a number of coding errors in several models in the `SimPyModels` folder.
- [FIX] Repaired `SimulationRT.py` bug introduced by recoding for the OO API.
- [FIX] Repaired errors in sample programs in documents:
  - Simulation with SimPy - In Depth Manual
  - SimPy's Object Oriented API Manual
  - Simulation With Real Time Synchronization Manual
  - SimPlot Manual
  - Publication-quality Plot Production With Matplotlib Manual

## 2.0.0 – 2009-01-26

This is a major release with changes to the SimPy application programming interface (API) and the formatting of the documentation.

### API changes

In addition to its existing API, SimPy now also has an object oriented API. The additional API

- allows running SimPy in parallel on multiple processors or multi-core CPUs,
- supports better structuring of SimPy programs,
- allows subclassing of class *Simulation* and thus provides users with the capability of creating new simulation modes/libraries like *SimulationTrace*, and
- reduces the total amount of SimPy code, thereby making it easier to maintain.

Note that the OO API is **in addition** to the old API. SimPy 2.0 is fully backward compatible.

### Documentation format changes

SimPy's documentation has been restructured and processed by the Sphinx documentation generation tool. This has generated one coherent, well structured document which can be easily browsed. A search capability is included.

## March 2008: Version 1.9.1

This is a bug-fix release which cures the following bugs:

- Excessive production of circular garbage, due to a circular reference between Process instances and event notices. This led to large memory requirements.
- Runtime error for preempts of processes holding multiple Resource objects.

It also adds a Short Manual, describing only the basic facilities of SimPy.

## December 2007: Version 1.9

This is a major release with added functionality/new user API calls and bug fixes.

### Major changes

- The event list handling has been changed to improve the runtime performance of large SimPy models (models with thousands of processes). The use of dictionaries for timestamps has been stopped. Thanks are due to Prof. Norm Matloff and a team of his students who did a study on improving SimPy performance. This was one of their recommendations. Thanks, Norm and guys! Furthermore, in version 1.9 the 'heapq' sorting package replaces 'bisect'. Finally, cancelling events no longer removes them, but rather marks them. When their event time comes, they are ignored. This was Tony Vignaux' idea!
- The Manual has been edited and given an easier-to-read layout.
- The Bank2 tutorial has been extended by models which use more advanced SimPy commands/constructs.

## Bug fixes

- The tracing of ‘activate’ statements has been enabled.

## Additions

- A method returning the time-weighted variance of observations has been added to classes Monitor and Tally.
- A shortcut activation method called “start” has been added to class Process.

## January 2007: Version 1.8

### Major Changes

- SimPy 1.8 and future releases will not run under the obsolete Python 2.2 version. They require Python 2.3 or later.
- The Manual has been thoroughly edited, restructured and rewritten. It is now also provided in PDF format.
- The Cheatsheet has been totally rewritten in a tabular format. It is provided in both XLS (MS Excel spreadsheet) and PDF format.
- The version of `SimPy.Simulation(RT/Trace/Step)` is now accessible by the variable ‘version’.
- The `__str__` method of Histogram was changed to return a table format.

### Bug fixes

- Repaired a bug in `yield waituntil` runtime code.
- Introduced check for `capacity` parameter of a Level or a Store being a number  $> 0$ .
- Added code so that `self.eventsFired` gets set correctly after an event fires in a compound yield `get/put` with a `waitevent` clause (reneging case).
- Repaired a bug in prettyprinting of Store objects.

### Additions

- New compound yield statements support time-out or event-based reneging in `get` and `put` operations on Store and Level instances.
- `yield get` on a Store instance can now have a filter function.
- All Monitor and Tally instances are automatically registered in list `allMonitors` and `allTallies`, respectively.
- The new function `startCollection` allows activation of Monitors and Tallies at a specified time.
- A `printHistogram` method was added to Tally and Monitor which generates a table-form histogram.
- In `SimPy.SimulationRT`: A function for allowing changing the ratio wall clock time to simulation time has been added.

## June 2006: Version 1.7.1

This is a maintenance release. The API has not been changed/added to.

- Repair of a bug in the `_get` methods of `Store` and `Level` which could lead to synchronization problems (blocking of producer processes, despite space being available in the buffer).
- Repair of `Level __init__` method to allow `initialBuffered` to be of either float or int type.
- Addition of type test for `Level get` parameter `'nrToGet'` to limit it to positive int or float.
- To improve pretty-printed output of `'Level'` objects, changed attribute `'_nrBuffered'` to `'nrBuffered'` (synonym for `'amount'` property).
- To improve pretty-printed output of `'Store'` objects, added attribute `'buffered'` (which refers to `'_theBuffer'` attribute).

## February 2006: Version 1.7

This is a major release.

- Addition of an abstract class `Buffer`, with two sub-classes `Store` and `Level`. `Store` and `Level` Buffers are used for modelling inter-process synchronization in producer/ consumer and multi-process cooperation scenarios.
- Addition of two new *yield* statements:
  - *yield put* for putting items into a buffer, and
  - *yield get* for getting items from a buffer.
- The Manual has undergone a major re-write/edit.
- All scripts have been restructured for compatibility with IronPython 1 beta2. This was done by moving all *import* statements to the beginning of the scripts. After the removal of the first (shebang) line, all scripts (with the exception of plotting and GUI scripts) can run successfully under this new Python implementation.

## September 2005: Version 1.6.1

This is a minor release.

- Addition of `Tally` data collection class as alternative to `Monitor`. It is intended for collecting very large data sets more efficiently in storage space and time than `Monitor`.
- Change of `Resource` to work with `Tally` (new `Resource` API is backwards-compatible with 1.6).
- Addition of function `setHistogram` to class `Monitor` for initializing histograms.
- New function `allEventNotices()` for debugging/teaching purposes. It returns a prettyprinted string with event times and names of process instances.
- Addition of function `allEventTimes` (returns event times of all scheduled events).

## 15 June 2005: Version 1.6

- Addition of two compound *yield* statement forms to support the modelling of processes reneging from resource queues.
- Addition of two test/demo files showing the use of the new reneging statements.
- Addition of test for prior simulation initialization in method `activate()`.

- Repair of bug in monitoring the waitQ of a resource when preemption occurs.
- Major restructuring/editing to Manual and Cheatsheet.

## 1 February 2005: Version 1.5.1

- MAJOR LICENSE CHANGE:

Starting with this version 1.5.1, SimPy is being released under the GNU Lesser General Public License (LGPL), instead of the GNU GPL. This change has been made to encourage commercial firms to use SimPy in for-profit work.

- Minor re-release
- No additional/changed functionality
- Includes unit test file 'MonitorTest.py' which had been accidentally deleted from 1.5
- Provides updated version of 'Bank.html' tutorial.
- Provides an additional tutorial ('Bank2.html') which shows how to use the new synchronization constructs introduced in SimPy 1.5.
- More logical, cleaner version numbering in files.

## 1 December 2004: Version 1.5

- No new functionality/API changes relative to 1.5 alpha
- Repaired bug related to waiting/queuing for multiple events
- SimulationRT: Improved synchronization with wallclock time on Unix/Linux

## 25 September 2004: Version 1.5alpha

- New functionality/API additions
  - SimEvents and signalling synchronization constructs, with 'yield waitevent' and 'yield queueevent' commands.
  - A general "wait until" synchronization construct, with the 'yield waituntil' command.
- No changes to 1.4.x API, i.e., existing code will work as before.

## 19 May 2004: Version 1.4.2

- Sub-release to repair two bugs:
  - The unittest for monitored Resource queues does not fail anymore.
  - SimulationTrace now works correctly with "yield hold,self" form.
- No functional or API changes

## 29 February 2004: Version 1.4.1

- Sub-release to repair two bugs:
  - The (optional) monitoring of the activeQ in Resource now works correctly.
  - The “cellphone.py” example is now implemented correctly.
- No functional or API changes

## 1 February 2004: Version 1.4

- Released on SourceForge.net

## 22 December 2003: Version 1.4 alpha

- New functionality/API changes
  - All classes in the SimPy API are now new style classes, i.e., they inherit from *object* either directly or indirectly.
  - Module *Monitor.py* has been merged into module *Simulation.py* and all *SimulationXXX.py* modules. Import of *Simulation* or any *SimulationXXX* module now also imports *Monitor*.
  - Some *Monitor* methods/attributes have changed. See Manual!
  - *Monitor* now inherits from *list*.
  - A class *Histogram* has been added to *Simulation.py* and all *SimulationXXX.py* modules.
  - A module *SimulationRT* has been added which allows synchronization between simulated and wallclock time.
  - A module *SimulationStep* which allows the execution of a simulation model event-by-event, with the facility to execute application code after each event.
  - A Tk/Tkinter-based module *SimGUI* has been added which provides a SimPy GUI framework.
  - A Tk/Tkinter-based module *SimPlot* has been added which provides for plot output from SimPy programs.

## 22 June 2003: Version 1.3

- No functional or API changes
- Reduction of sourcecode linelength in *Simulation.py* to  $\leq 80$  characters

## June 2003: Version 1.3 alpha

- Significantly improved performance
- Significant increase in number of quasi-parallel processes SimPy can handle
- New functionality/API changes:
  - Addition of *SimulationTrace*, an event trace utility
  - Addition of *Lister*, a prettyprinter for instance attributes

- No API changes
- Internal changes:
  - Implementation of a proposal by Simon Frost: storing the keys of the event set dictionary in a binary search tree using bisect. Thank you, Simon! SimPy 1.3 is dedicated to you!
- Update of Manual to address tracing.
- Update of Interfacing doc to address output visualization using Scientific Python gplt package.

## 29 April 2003: Version 1.2

- No changes in API.
- **Internal changes:**
  - Defined “True” and “False” in Simulation.py to support Python 2.2.

## 22 October 2002

- Re-release of 0.5 Beta on SourceForge.net to replace corrupted file `__init__.py`.
- No code changes whatever!

## 18 October 2002

- Version 0.5 Beta-release, intended to get testing by application developers and system integrators in preparation of first full (production) release. Released on SourceForge.net on 20 October 2002.
- More models
- Documentation enhanced by a manual, a tutorial (“The Bank”) and installation instructions.
- Major changes to the API:
  - Introduced ‘simulate(until=0)’ instead of ‘scheduler(till=0)’. Left ‘scheduler()’ in for backward compatibility, but marked as deprecated.
  - Added attribute “name” to class Process. Process constructor is now:

```
def __init__(self, name="a_process")
```

Backward compatible if keyword parameters used.

- Changed Resource constructor to:

```
def __init__(self, capacity=1, name="a_resource", unitName="units")
```

Backward compatible if keyword parameters used.

## 27 September 2002

- Version 0.2 Alpha-release, intended to attract feedback from users
- Extended list of models
- Updated documentation

## 17 September 2002

- Version 0.1.2 published on SourceForge; fully working, pre-alpha code
- Implements simulation, shared resources with queuing (FIFO), and monitors for data gathering/analysis.
- Contains basic documentation (cheatsheet) and simulation models for test and demonstration.

## Acknowledgments

SimPy 2 has been primarily developed by Stefan Scherfke and Ontje Lünsdorf, starting from SimPy 1.9. Their work has resulted in a most elegant combination of an object oriented API with the existing API, maintaining full backward compatibility. It has been quite easy to integrate their product into the existing SimPy code and documentation environment.

Thanks, guys, for this great job! **SimPy 2.0 is dedicated to you!**

SimPy was originally created by Klaus Müller and Tony Vignaux. They pushed its development for several years and built the SimPy community. Without them, there would be no SimPy 3.

Thanks, guys, for this great job! **SimPy 3.0 is dedicated to you!**

The many contributions of the SimPy user and developer communities are of course also gratefully acknowledged.

## Ports and comparable libraries

Reimplementations of SimPy and libraries similar to SimPy are available in the following languages:

- C#: [SimSharp](#) (written by Andreas Beham)
- Julia: [SimJulia](#)
- R: [Simmer](#)

## Defense of Design

This document explains why SimPy is designed the way it is and how its design evolved over time.

### Original Design of SimPy 1

SimPy 1 was heavily inspired by *Simula 67* and *Simscrip*t. The basic entity of the framework was a process. A process described a temporal sequence of actions.

In SimPy 1, you implemented a process by sub-classing `Process`. The instance of such a subclass carried both, process and simulation internal information, whereas the latter wasn't of any use to the process itself. The sequence of actions of the process was specified in a method of the subclass, called the *process execution method* (or PEM in short). A PEM interacted with the simulation by yielding one of several keywords defined in the simulation package.

The simulation itself was executed via module level functions. The simulation state was stored in the global scope. This made it very easy to implement and execute a simulation (despite from having to inherit from `Process` and instantiating the processes before starting their PEMs). However, having all simulation state global makes it hard to parallelize multiple simulations.

SimPy 1 also followed the “batteries included” approach, providing shared resources, monitoring, plotting, GUIs and multiple types of simulations (“normal”, real-time, manual stepping, with tracing).

The following code fragment shows how a simple simulation could be implemented in SimPy 1:

```
from SimPy.Simulation import Process, hold, initialize, activate, simulate

class MyProcess(Process):
    def pem(self, repeat):
        for i in range(repeat):
            yield hold, self, 1

initialize()
proc = MyProcess()
activate(proc, proc.pem(3))
simulate(until=10)

sim = Simulation()
proc = MyProcess(sim=sim)
sim.activate(proc, proc.pem(3))
sim.simulate(until=10)
```

## Changes in SimPy 2

SimPy 2 mostly stuck with SimPy 1’s design, but added an object orient API for the execution of simulations, allowing them to be executed in parallel. Since processes and the simulation state were so closely coupled, you now needed to pass the `Simulation` instance into your process to “bind” them to that instance. Additionally, you still had to activate the process. If you forgot to pass the simulation instance, the process would use a global instance thereby breaking your program. SimPy 2’s OO-API looked like this:

```
from SimPy.Simulation import Simulation, Process, hold

class MyProcess(Process):
    def pem(self, repeat):
        for i in range(repeat):
            yield hold, self, 1

sim = Simulation()
proc = MyProcess(sim=sim)
sim.activate(proc, proc.pem(3))
sim.simulate(until=10)
```

## Changes and Decisions in SimPy 3

The original goals for SimPy 3 were to simplify and PEP8-ify its API and to clean up and modularize its internals. We knew from the beginning that our goals would not be achievable without breaking backwards compatibility with SimPy 2. However, we didn’t expect the API changes to become as extensive as they ended up to be.

We also removed some of the included batteries, namely SimPy’s plotting and GUI capabilities, since dedicated libraries like `matplotlib` or `PySide` do a much better job here.

However, by far the most changes are—from the end user’s view—mostly syntactical. Thus, porting from 2 to 3 usually just means replacing a line of SimPy 2 code with its SimPy3 equivalent (e.g., replacing `yield hold, self, 1` with `yield env.timeout(1)`).

In short, the most notable changes in SimPy 3 are:

- No more sub-classing of `Process` required. PEMs can even be simple module level functions.
- The simulation state is now stored in an `Environment` which can also be used by a PEM to interact with the simulation.
- PEMs now yield event objects. This implicates interesting new features and allows an easy extension with new event types.

These changes are causing the above example to now look like this:

```
from simpy import Environment, simulate

def pem(env, repeat):
    for i in range(repeat):
        yield env.timeout(i)

env = Environment()
env.process(pem(env, 7))
simulate(env, until=10)
```

The following sections describe these changes in detail:

### No More Sub-classing of `Process`

In SimPy 3, every Python generator can be used as a PEM, no matter if it is a module level function or a method of an object. This reduces the amount of code required for simple processes. The `Process` class still exists, but you don't need to instantiate it by yourself, though. More on that later.

### Processes Live in an Environment

Process and simulation state are decoupled. An `Environment` holds the simulation state and serves as base API for processes to create new events. This allows you to implement advanced use cases by extending the `Process` or `Environment` class without affecting other components.

For the same reason, the `simulate()` method now is a module level function that takes an environment to simulate.

### Stronger Focus on Events

In former versions, PEMs needed to yield one of SimPy's built-in keywords (like `hold`) to interact with the simulation. These keywords had to be imported separately and were bound to some internal functions that were tightly integrated with the `Simulation` and `Process` making it very hard to extend SimPy with new functionality.

In SimPy 3, PEMs just need to yield events. There are various built-in event types, but you can also create custom ones by making a subclass of a `BaseEvent`. Most events are generated by factory methods of `Environment`. For example, `Environment.timeout()` creates a `Timeout` event that replaces the `hold` keyword.

The `Process` is now also an event. You can now yield another process and wait for it to finish. For example, think of a car-wash simulation where "washing" is a process that the car processes can wait for once they enter the washing station.

## Creating Events via the Environment or Resources

The `Environment` and resources have methods to create new events, e.g. `Environment.timeout()` or `Resource.request()`. Each of these methods maps to a certain event type. It creates a new instance of it and returns it, e.g.:

```
def event(self):
    return Event()
```

To simplify things, we wanted to use the event classes directly as methods:

```
class Environment(object):
    event = Event
```

This was, unfortunately, not directly possible and we had to wrap the classes to behave like bound methods. Therefore, we introduced a `BoundClass`:

```
class BoundClass(object):
    """Allows classes to behave like methods. The ``__get__()`` descriptor
    is basically identical to ``function.__get__()`` and binds the first
    argument of the ``cls`` to the descriptor instance.

    """
    def __init__(self, cls):
        self.cls = cls

    def __get__(self, obj, type=None):
        if obj is None:
            return self.cls
        return types.MethodType(self.cls, obj)

class Environment(object):
    event = BoundClass(Event)
```

These methods are called a lot, so we added the event classes as `types.MethodType` to the instance of `Environment` (or the resources, respectively):

```
class Environment(object):
    def __init__(self):
        self.event = types.MethodType(Event, self)
```

It turned out the the class attributes (the `BoundClass` instances) were now quite useless, so we removed them although it was actually the “right” way to add classes as methods to another class.

## Release Process

This process describes the steps to execute in order to release a new version of SimPy.

### Preparations

1. Close all tickets for the next version.
2. Update the *minium* required versions of dependencies in `setup.py`. Update the *exact* version of all entries in `requirements.txt`.

3. Run `tox` from the project root. All tests for all supported versions must pass:

```
$ tox
[...]
_____ summary _____
py27: commands succeeded
py32: commands succeeded
py33: commands succeeded
pypy: commands succeeded
congratulations :)
```

**Nota:** Tox will use the `requirements.txt` to setup the venvs, so make sure you've updated it!

4. Build the docs (HTML is enough). Make sure there are no errors and undefined references.

```
$ cd docs/
$ make clean html
$ cd ..
```

5. Check if all authors are listed in `AUTHORS.txt`.
6. Update the change logs (`CHANGES.txt` and `docs/about/history.rst`). Only keep changes for the current major release in `CHANGES.txt` and reference the history page from there.
7. Commit all changes:

```
$ hg ci -m 'Updated change log for the upcoming release.'
```

8. Update the version number in `simpy/__init__.py` and commit:

```
$ hg ci -m 'Bump version from x.y.z to a.b.c'
```

**Aviso:** Do not yet tag and push the changes so that you can safely do a rollback if one of the next step fails and you need change something!

9. Write a draft for the announcement mail with a list of changes, acknowledgements and installation instructions. Everyone in the team should agree with it.

## Build and release

1. Test the release process. Build a source distribution and a `wheel` package and test them:

```
$ python setup.py sdist bdist_wheel
$ ls dist/
simpy-a.b.c-py2.py3-none-any.whl simpy-a.b.c.tar.gz
```

Try installing them:

```
$ rm -rf /tmp/simpy-sdist # ensure clean state if ran repeatedly
$ virtualenv /tmp/simpy-sdist
$ /tmp/simpy-sdist/bin/pip install pytest
$ /tmp/simpy-sdist/bin/pip install dist/simpy-a.b.c.tar.gz
$ /tmp/simpy-sdist/bin/python
>>> import simpy # doctest: +SKIP
>>> simpy.__version__ # doctest: +SKIP
'a.b.c'
>>> simpy.test() # doctest: +SKIP
```

and

```
$ rm -rf /tmp/simpy-wheel # ensure clean state if ran repeatedly
$ virtualenv /tmp/simpy-wheel
$ /tmp/simpy-wheel/bin/pip install pytest
$ /tmp/simpy-wheel/bin/pip install dist/simpy-a.b.c-py2.py3-none-any.whl
$ /tmp/simpy-wheel/bin/python
>>> import simpy # doctest: +SKIP
>>> simpy.__version__ # doctest: +SKIP
'a.b.c'
>>> simpy.test() # doctest: +SKIP
```

2. Create or check your accounts for the *test server* <<https://testpypi.python.org/pypi>> and PyPI. Update your `~/.pypirc` with your current credentials:

```
[distutils]
index-servers =
    pypi
    test

[test]
repository = https://testpypi.python.org/pypi
username = <your test user name goes here>
password = <your test password goes here>

[pypi]
repository = http://pypi.python.org/pypi
username = <your production user name goes here>
password = <your production password goes here>
```

3. Upload the distributions for the new version to the test server and test the installation again:

```
$ twine upload -r test dist/simpy*a.b.c*
$ pip install -i https://testpypi.python.org/pypi simpy
```

4. Check if the package is displayed correctly: <https://testpypi.python.org/pypi/simpy>
5. Finally upload the package to PyPI and test its installation one last time:

```
$ twine upload -r pypi dist/simpy*a.b.c*
$ pip install -U simpy
```

6. Check if the package is displayed correctly: <https://pypi.python.org/pypi/simpy>

## Post release

1. Push your changes:

```
$ hg tag a.b.c
$ hg push ssh://hg@bitbucket.org/simpy/simpy
```

2. Activate the [documentation build](#) for the new version.
3. Send the prepared email to the mailing list and post it on Google+.
4. Update [Wikipedia](#) entries.
5. Update [Python Wiki](#)
6. Post something to Planet Python (e.g., via Stefan's blog).

## License

The MIT License (MIT)

Copyright (c) 2013 Ontje Lünsdorf and Stefan Scherfke (also see AUTHORS.txt)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

---

## Indices and tables

---

- `genindex`
- `search`

## S

- `simpy`, 61
- `simpy.core`, 62
- `simpy.events`, 64
- `simpy.exceptions`, 64
- `simpy.resources`, 72
- `simpy.resources.base`, 77
- `simpy.resources.container`, 75
- `simpy.resources.resource`, 73
- `simpy.resources.store`, 76
- `simpy.rt`, 79
- `simpy.util`, 80

**A**

active\_process (atributo `simpy.core.BaseEnvironment`), 62  
 active\_process (atributo `simpy.core.Environment`), 63  
 active\_process (atributo `simpy.rt.RealtimeEnvironment`), 79  
 all\_events() (método estático `simpy.events.Condition`), 70  
 all\_of() (método `simpy.core.Environment`), 63  
 all\_of() (método `simpy.rt.RealtimeEnvironment`), 79  
 AllOf (classe em `simpy.events`), 71  
 amount (atributo `simpy.resources.container.ContainerGet`), 75  
 amount (atributo `simpy.resources.container.ContainerPut`), 75  
 any\_events() (método estático `simpy.events.Condition`), 70  
 any\_of() (método `simpy.core.Environment`), 63  
 any\_of() (método `simpy.rt.RealtimeEnvironment`), 79  
 AnyOf (classe em `simpy.events`), 71  
 append() (método `simpy.resources.resource.SortedQueue`), 75

**B**

BaseEnvironment (classe em `simpy.core`), 62  
 BaseResource (classe em `simpy.resources.base`), 77  
 bind\_early() (método estático `simpy.core.BoundClass`), 63  
 BoundClass (classe em `simpy.core`), 63  
 by (atributo `simpy.resources.resource.Preempted`), 74

**C**

callbacks (atributo `simpy.events.Event`), 65  
 cancel() (método `simpy.resources.base.Get`), 78  
 cancel() (método `simpy.resources.base.Put`), 78  
 capacity (atributo `simpy.resources.base.BaseResource`), 78  
 cause (atributo `simpy.exceptions.Interrupt`), 64  
 Condition (classe em `simpy.events`), 69  
 ConditionValue (classe em `simpy.events`), 72  
 Container (classe em `simpy.resources.container`), 75

ContainerGet (classe em `simpy.resources.container`), 75  
 ContainerPut (classe em `simpy.resources.container`), 75  
 count (atributo `simpy.resources.resource.Resource`), 73

**D**

defused (atributo `simpy.events.AllOf`), 71  
 defused (atributo `simpy.events.AnyOf`), 71  
 defused (atributo `simpy.events.Condition`), 70  
 defused (atributo `simpy.events.Event`), 65  
 defused (atributo `simpy.events.Initialize`), 67  
 defused (atributo `simpy.events.Interruption`), 67  
 defused (atributo `simpy.events.Process`), 69  
 defused (atributo `simpy.events.Timeout`), 66

**E**

EmptySchedule (classe em `simpy.core`), 64  
 env (atributo `simpy.events.Event`), 65  
 Environment (classe em `simpy.core`), 62  
 Event (classe em `simpy.events`), 65  
 event() (método `simpy.core.Environment`), 63  
 event() (método `simpy.rt.RealtimeEnvironment`), 79  
 exit() (método `simpy.core.BaseEnvironment`), 62  
 exit() (método `simpy.core.Environment`), 63  
 exit() (método `simpy.rt.RealtimeEnvironment`), 79

**F**

factor (atributo `simpy.rt.RealtimeEnvironment`), 79  
 fail() (método `simpy.events.AllOf`), 71  
 fail() (método `simpy.events.AnyOf`), 72  
 fail() (método `simpy.events.Condition`), 70  
 fail() (método `simpy.events.Event`), 66  
 fail() (método `simpy.events.Initialize`), 67  
 fail() (método `simpy.events.Interruption`), 68  
 fail() (método `simpy.events.Process`), 69  
 fail() (método `simpy.events.Timeout`), 66  
 filter (atributo `simpy.resources.store.FilterStoreGet`), 77  
 FilterStore (classe em `simpy.resources.store`), 76  
 FilterStoreGet (classe em `simpy.resources.store`), 77

**G**

get (atributo `simpy.resources.base.BaseResource`), 78

- get (atributo `simpy.resources.container.Container`), 75  
get (atributo `simpy.resources.store.FilterStore`), 77  
get (atributo `simpy.resources.store.Store`), 76  
Get (classe em `simpy.resources.base`), 78  
get\_queue (atributo `simpy.resources.base.BaseResource`), 78  
GetQueue (atributo `simpy.resources.base.BaseResource`), 77  
GetQueue (atributo `simpy.resources.resource.PriorityResource`), 73
- I**
- Infinity (no módulo `simpy.core`), 64  
Initialize (classe em `simpy.events`), 67  
Interrupt, 64  
interrupt() (método `simpy.events.Process`), 69  
Interruption (classe em `simpy.events`), 67  
is\_alive (atributo `simpy.events.Process`), 69  
item (atributo `simpy.resources.store.PriorityItem`), 76  
item (atributo `simpy.resources.store.StorePut`), 77  
items (atributo `simpy.resources.store.Store`), 76
- K**
- key (atributo `simpy.resources.resource.PriorityRequest`), 74
- L**
- level (atributo `simpy.resources.container.Container`), 75
- M**
- maxlen (atributo `simpy.resources.resource.SortedQueue`), 75
- N**
- NORMAL (no módulo `simpy.events`), 64  
now (atributo `simpy.core.BaseEnvironment`), 62  
now (atributo `simpy.core.Environment`), 63  
now (atributo `simpy.rt.RealtimeEnvironment`), 79
- O**
- ok (atributo `simpy.events.AllOf`), 71  
ok (atributo `simpy.events.AnyOf`), 72  
ok (atributo `simpy.events.Condition`), 70  
ok (atributo `simpy.events.Event`), 65  
ok (atributo `simpy.events.Initialize`), 67  
ok (atributo `simpy.events.Interruption`), 68  
ok (atributo `simpy.events.Process`), 69  
ok (atributo `simpy.events.Timeout`), 66
- P**
- peek() (método `simpy.core.Environment`), 63  
peek() (método `simpy.rt.RealtimeEnvironment`), 79  
PENDING (no módulo `simpy.events`), 64  
preempt (atributo `simpy.resources.resource.PriorityRequest`), 74  
Preempted (classe em `simpy.resources.resource`), 74  
PreemptiveResource (classe em `simpy.resources.resource`), 74  
priority (atributo `simpy.resources.resource.PriorityRequest`), 74  
priority (atributo `simpy.resources.store.PriorityItem`), 76  
PriorityItem (classe em `simpy.resources.store`), 76  
PriorityRequest (classe em `simpy.resources.resource`), 74  
PriorityResource (classe em `simpy.resources.resource`), 73  
PriorityStore (classe em `simpy.resources.store`), 76  
Process (classe em `simpy.events`), 68  
process() (método `simpy.core.Environment`), 63  
process() (método `simpy.rt.RealtimeEnvironment`), 79  
processed (atributo `simpy.events.AllOf`), 71  
processed (atributo `simpy.events.AnyOf`), 72  
processed (atributo `simpy.events.Condition`), 70  
processed (atributo `simpy.events.Event`), 65  
processed (atributo `simpy.events.Initialize`), 67  
processed (atributo `simpy.events.Interruption`), 68  
processed (atributo `simpy.events.Process`), 69  
processed (atributo `simpy.events.Timeout`), 66  
put (atributo `simpy.resources.base.BaseResource`), 78  
put (atributo `simpy.resources.container.Container`), 75  
put (atributo `simpy.resources.store.FilterStore`), 77  
put (atributo `simpy.resources.store.Store`), 76  
Put (classe em `simpy.resources.base`), 78  
put\_queue (atributo `simpy.resources.base.BaseResource`), 78  
PutQueue (atributo `simpy.resources.base.BaseResource`), 77  
PutQueue (atributo `simpy.resources.resource.PriorityResource`), 73
- Q**
- queue (atributo `simpy.resources.resource.Resource`), 73
- R**
- RealtimeEnvironment (classe em `simpy.rt`), 79  
release (atributo `simpy.resources.resource.PriorityResource`), 74  
release (atributo `simpy.resources.resource.Resource`), 73  
Release (classe em `simpy.resources.resource`), 74  
request (atributo `simpy.resources.resource.PriorityResource`), 73  
request (atributo `simpy.resources.resource.Release`), 75  
request (atributo `simpy.resources.resource.Resource`), 73  
Request (classe em `simpy.resources.resource`), 74  
resource (atributo `simpy.resources.resource.Preempted`), 74  
Resource (classe em `simpy.resources.resource`), 73  
run() (método `simpy.core.BaseEnvironment`), 62

run() (método `simpy.core.Environment`), 63  
 run() (método `simpy.rt.RealtimeEnvironment`), 80

## S

schedule() (método `simpy.core.BaseEnvironment`), 62  
 schedule() (método `simpy.core.Environment`), 63  
 schedule() (método `simpy.rt.RealtimeEnvironment`), 79  
 simpy (módulo), 61  
 simpy.core (módulo), 62  
 simpy.events (módulo), 64  
 simpy.exceptions (módulo), 64  
 simpy.resources (módulo), 72  
 simpy.resources.base (módulo), 77  
 simpy.resources.container (módulo), 75  
 simpy.resources.resource (módulo), 73  
 simpy.resources.store (módulo), 76  
 simpy.rt (módulo), 79  
 simpy.util (módulo), 80  
 SimPyException, 64  
 SortedQueue (classe em `simpy.resources.resource`), 75  
 start\_delayed() (no módulo `simpy.util`), 80  
 step() (método `simpy.core.BaseEnvironment`), 62  
 step() (método `simpy.core.Environment`), 63  
 step() (método `simpy.rt.RealtimeEnvironment`), 79  
 StopProcess, 64  
 Store (classe em `simpy.resources.store`), 76  
 StoreGet (classe em `simpy.resources.store`), 77  
 StorePut (classe em `simpy.resources.store`), 77  
 strict (atributo `simpy.rt.RealtimeEnvironment`), 79  
 succeed() (método `simpy.events.AllOf`), 71  
 succeed() (método `simpy.events.AnyOf`), 72  
 succeed() (método `simpy.events.Condition`), 70  
 succeed() (método `simpy.events.Event`), 66  
 succeed() (método `simpy.events.Initialize`), 67  
 succeed() (método `simpy.events.Interruption`), 68  
 succeed() (método `simpy.events.Process`), 69  
 succeed() (método `simpy.events.Timeout`), 66  
 sync() (método `simpy.rt.RealtimeEnvironment`), 79

## T

target (atributo `simpy.events.Process`), 68  
 time (atributo `simpy.resources.resource.PriorityRequest`), 74  
 Timeout (classe em `simpy.events`), 66  
 timeout() (método `simpy.core.Environment`), 63  
 timeout() (método `simpy.rt.RealtimeEnvironment`), 79  
 trigger() (método `simpy.events.AllOf`), 71  
 trigger() (método `simpy.events.AnyOf`), 72  
 trigger() (método `simpy.events.Condition`), 70  
 trigger() (método `simpy.events.Event`), 65  
 trigger() (método `simpy.events.Initialize`), 67  
 trigger() (método `simpy.events.Interruption`), 68  
 trigger() (método `simpy.events.Process`), 69  
 trigger() (método `simpy.events.Timeout`), 66

triggered (atributo `simpy.events.AllOf`), 71  
 triggered (atributo `simpy.events.AnyOf`), 72  
 triggered (atributo `simpy.events.Condition`), 70  
 triggered (atributo `simpy.events.Event`), 65  
 triggered (atributo `simpy.events.Initialize`), 67  
 triggered (atributo `simpy.events.Interruption`), 68  
 triggered (atributo `simpy.events.Process`), 69  
 triggered (atributo `simpy.events.Timeout`), 66

## U

URGENT (no módulo `simpy.events`), 64  
 usage\_since (atributo `simpy.resources.resource.Preempted`), 74  
 usage\_since (atributo `simpy.resources.resource.PriorityRequest`), 74  
 users (atributo `simpy.resources.resource.Resource`), 73

## V

value (atributo `simpy.events.AllOf`), 71  
 value (atributo `simpy.events.AnyOf`), 72  
 value (atributo `simpy.events.Condition`), 70  
 value (atributo `simpy.events.Event`), 65  
 value (atributo `simpy.events.Initialize`), 67  
 value (atributo `simpy.events.Interruption`), 68  
 value (atributo `simpy.events.Process`), 69  
 value (atributo `simpy.events.Timeout`), 66  
 value (atributo `simpy.exceptions.StopProcess`), 64