

For example, suppose that the contract's code is:

```
if !contract.storage[msg.data[0]]:  
    contract.storage[msg.data[0]] = msg.data[1]
```

Note that in reality the contract code is written in the low-level EVM code; this example is written in Serpent, our high-level language, for clarity, and can be compiled down to EVM code. Suppose that the contract's storage starts off empty, and a transaction is sent with 10 ether value, 2000 gas, 0.001 ether gasprice, and two data fields: [2, 'CHARLIE']^[3]. The process for the state transition function in this case is as follows:

1. Check that the transaction is valid and well formed.
2. Check that the transaction sender has at least $2000 * 0.001 = 2$ ether. If it is, then subtract 2 ether from the sender's account.
3. Initialize gas = 2000; assuming the transaction is 170 bytes long and the byte-fee is 5, subtract 850 so that there is 1150 gas left.
4. Subtract 10 more ether from the sender's account, and add it to the contract's account.
5. Run the code. In this case, this is simple: it checks if the contract's storage at index 2 is used, notices that it is not, and so it sets the storage at index 2 to the value CHARLIE. Suppose this takes 187 gas, so the remaining amount of gas is $1150 - 187 = 963$
6. Add $963 * 0.001 = 0.963$ ether back to the sender's account, and return the resulting state.

If there was no contract at the receiving end of the transaction, then the total transaction fee would simply be equal to the provided GASPRICE multiplied by the length of the transaction in bytes, and the data sent alongside the transaction would be irrelevant. Additionally, note that contract-initiated messages can assign a gas limit to the computation that they spawn, and if the sub-computation runs out of gas it gets reverted only to the point of the message call. Hence, just like transactions, contracts can secure their limited computational resources by setting strict limits on the sub-computations that they spawn.

