

11 Circom

11.1 Circom Walkthrough

In this final lecture, we bridge the gap between the theoretical concepts presented in previous lectures and their practical realization using the **Circom** DSL.

Definition 11.1. **Circom** is a domain-specific language for building arithmetic circuits that can be used to produce zk-SNARK proofs.

Throughout this lecture, we will walk through how concepts like R1CS, witness, trusted setup, and verification keys appear in actual code and practice.

11.1.1 Journey Begins

In the previous lectures, we covered a variety of theoretical concepts: zk-SNARKs, trusted setup, arithmetic circuits, constraints, witnesses, and the Rank-1 Constraint System (R1CS) representation. Now, let's see how these appear in practice.

11.1.2 From Theory to Practice: Circom Basics

We learned that a circuit can represent a complex arithmetic computation over a finite field. Circom allows us to write these circuits in a high-level syntax. To begin, consider the arithmetic circuit $r = x \times y$.

It can be represented in Circom syntax as follows:

```

1  pragma circom 2.1.6;
2
3  template Math() {
4      signal input x;
5      signal input y;
6
7      signal output r <== x * y;
8  }
9
10 component main = Math();

```

Here, we see how easy it is to define a circuit that takes two inputs x, y and outputs their product r . The `template` defines a reusable circuit component, while `signal input` and `signal output` represent inputs and outputs, respectively. Intermediate signals (without input or output) are internal primitives within the circuit.

Public vs Private Signals:

Output signals are always public. You may also define public inputs by specifying them in the main component, for example:

```
component main {public [x]} = Math();
```

This means x is a public input and will appear in the verification context. The order of public signals in the final proof verification step follows the order of their definition inside the template, starting with outputs.

For example, if your circuit looks like this:

```
1     template Circuit() {
2         signal x;
3
4         signal output o2;
5
6         signal input c;
7         signal input a;
8
9         signal k1;
10
11        signal input b;
12
13        signal output o1;
14    }
15
16    component main {public [a, b, c]} = Circuit();
```

The order of the public signals that should be passed to the verifier is as follows:

(o2, o1, c, a, b)

11.1.3 Arguments, Functions, and Vars

Sometimes we need to calculate some values as constants for our circuit. For example, if you want your circuit to be a multi-tool that, based on provided arguments, can work with different cases. For this purpose, we can declare *functions* and *vars* inside the circuit, as shown below:

```
1     function transformNumber(value) {
2         return value ** 2;
3     }
4
5     template Math(padding) {
6         signal input x;
7         signal input y;
8
9         var elementsNumber = transformNumber(padding);
10        signal b <== x * elementsNumber;
11
12        signal output r <== b * y;
13    }
14
15    component main {public [x]} = Math(12);
```

Here, `var elementsNumber` and the function `transformNumber` are evaluated at compile time. Remember that assignments using `var` and functions do not produce constraints by themselves. Only `<==`, `==>`, or `===` and actual arithmetic on signals produce constraints reflected in R1CS.

Remark. Sometimes, one needs to perform operations like division or non-quadratic multiplication on the signal. For this purpose, you can use the `-->` and `<--` notations to compute the values “out-of-circuit”. For example:

```
template Math() {
    signal input x;
    signal input y;

    signal b <-- x / y;

    signal output r <== b * y;
}

component main = Math();
```

In this case, no constraints are generated with the x input, and it does not even participate in the witness directly.

Notice! This is the main difference between Circom and other languages. When you write a function evaluation $y = f(x)$ in any other language (say, Python or Rust), you are specifying the set of instructions to compute y from x (commonly sequentially). In Circom (or in any other R1CS language) you are merely asserting the **correctness** of the computation and therefore of all intermediate computations. This way, if your task would have been to compute $y = \frac{1}{x}$, you could simply ask y to be the result of the division (that can be computed out of circuit) and then asserting that $x \times y = 1$ with $x \neq 0$. This way, you are not writing the division itself, but the constraints that the division should satisfy.

11.1.4 Theoretical Recap: Using the Learned Concepts

Next, let us apply the learned concepts to a more complex examples. We start with the `if` statement logic.

Example. Recall the complex example we analyzed in earlier lectures:

```
def r(x1: bool, x2: F, x3: F) -> F:
    return x2 * x3 if x1 else x2 + x3
```

This can be represented as:

$$r = x_1 \times (x_2 \times x_3) + (1 - x_1) \times (x_2 + x_3).$$

We also had the additional constraint $x_1 \times (1 - x_1) = 0$ to ensure x_1 is binary. The resulting system of constraints was:

$$x_1 \times x_1 = x_1 \tag{1}$$

$$x_2 \times x_3 = \text{mult} \tag{2}$$

$$x_1 \times \text{mult} = \text{selectMult} \tag{3}$$

$$(1 - x_1) \times (x_2 + x_3) = r - \text{selectMult} \tag{4}$$

It took us quite some time to understand and come up with the constraint system, which can be visualized as follows:




Figure 11.1: Example of a circuit evaluating the if statement logic.

The inputs can be directly transformed into signals like below:

```

1  template Math() {
2      signal output r;
3
4      signal input x1;
5
6      signal input x2;
7      signal input x3;
8  }
```

In our case, we have an additional output signal, so we can "return" it from the circuit. Now, let's compare the mathematical and Circom representations.

Mathematical Constraints:

$$\begin{aligned}
 x_1 \times x_1 &= x_1 \\
 x_2 \times x_3 &= \text{mult} \\
 x_1 \times \text{mult} &= \text{selectMult} \\
 (1 - x_1) \times (x_2 + x_3) &= r - \text{selectMult}
 \end{aligned}$$

Circom Representation:

```

x1 * x1 === x1;
signal mult <== x2 * x3;
signal selectMult <== x1 * mult;
(1 - x1) * (x2 + x3) + selectMult ==> r;
```

As we can see, the translation from math to Circom is straightforward. We have used *signals* for constraint definitions.

Remark. If you wish to follow along with the explanations in the following chapters:

1. Clone the repository <https://github.com/ZKDL-Camp/hardhat-zkit-template>.
2. Run `npm install` to install dependencies and `npx hardhat zkit make` to compile Circom circuits and generate the necessary artifacts.

11.1.5 From R1CS to Proof Generation

Now, let us break down everything that is happening during the proof generation process. After compilation, you will find the following files in the `zkit/artifacts/circuits` folder

(starting from the project root):

- `.r1cs` file: The Rank-1 Constraint System representation of the circuit.
- `.wasm` and `*.js` files: The code to compute the witness from the given inputs.
- `.zkey` file: Proving keys after the trusted setup.
- `.sym` file: Symbolic reference for signals.

R1CS File. Let us start with the `.r1cs` file. In Section 8, we defined the following coefficient vectors (in simple constraints) for our task:

$$\begin{array}{lll}
 \mathbf{a}_1 = (0, 0, 1, 0, 0, 0, 0) & \mathbf{b}_1 = (0, 0, 1, 0, 0, 0, 0) & \mathbf{c}_1 = (0, 0, 1, 0, 0, 0, 0) \\
 \mathbf{a}_2 = (0, 0, 0, 1, 0, 0, 0) & \mathbf{b}_2 = (0, 0, 0, 0, 1, 0, 0) & \mathbf{c}_2 = (0, 0, 0, 0, 0, 1, 0) \\
 \mathbf{a}_3 = (0, 0, 1, 0, 0, 0, 0) & \mathbf{b}_3 = (0, 0, 0, 0, 0, 1, 0) & \mathbf{c}_3 = (0, 0, 0, 0, 0, 0, 1) \\
 \mathbf{a}_4 = (1, 0, -1, 0, 0, 0, 0) & \mathbf{b}_4 = (0, 0, 0, 1, 1, 0, 0) & \mathbf{c}_4 = (0, 1, 0, 0, 0, 0, -1)
 \end{array}$$

On the other hand, using the test from the `test/Math.witness.test.ts` file and reading the R1CS file, we can see:

test/Math.witness.test.ts

```

expect(constraint1[0]).to.deep.equal([ 0n, 0n, 1n, 0n, 0n, 0n, 0n ]);
expect(constraint1[1]).to.deep.equal([ 0n, 0n, 1n, 0n, 0n, 0n, 0n ]);
expect(constraint1[2]).to.deep.equal([ 0n, 0n, 1n, 0n, 0n, 0n, 0n ]);

expect(constraint2[0]).to.deep.equal([ 0n, 0n, 0n, babyJub.F.negone, 0n, 0n, 0n ]);
expect(constraint2[1]).to.deep.equal([ 0n, 0n, 0n, 0n, 1n, 0n, 0n ]);
expect(constraint2[2]).to.deep.equal([ 0n, 0n, 0n, 0n, 0n, babyJub.F.negone, 0n ]);

expect(constraint3[0]).to.deep.equal([ 0n, 0n, babyJub.F.negone, 0n, 0n, 0n, 0n ]);
expect(constraint3[1]).to.deep.equal([ 0n, 0n, 0n, 0n, 0n, 1n, 0n ]);
expect(constraint3[2]).to.deep.equal([ 0n, 0n, 0n, 0n, 0n, 0n, babyJub.F.negone ]);

expect(constraint4[0]).to.deep.equal([ babyJub.F.negone, 0n, 0n, 0n, 0n, 0n, 0n ]);
expect(constraint4[1]).to.deep.equal([ 0n, 0n, 0n, 1n, 0n, 0n, 0n ]);
expect(constraint4[2]).to.deep.equal([ 0n, babyJub.F.negone, 0n, 0n, 0n, 0n, 0n ]);

```

Mostly, the structure generated by Circom aligns with what we had devised, except for the last constraint. The difference occurs because of Circom's optimization to make proof generation and verification more efficient.

Now, let's take a closer look at how the witness is computed. In Section 8, we had:

$$\mathbf{w} = (1, r, x_1, x_2, x_3, \text{mult}, \text{selectMult})$$

Given the inputs: $x_1 = 1, x_2 = 3, x_3 = 4$, we can quickly do the math and find out that the actual witness should look like this: $\mathbf{w} = (1, 12, 1, 3, 4, 12, 12)$ based on:

$$\begin{aligned}
 \text{mult} &= 3 \times 4 = 12 \\
 \text{selectMult} &= 1 \times 12 = 12 \\
 r &= 1 \times (3 \times 4) + (1 - 1) \times (3 + 4) = 12 + 0 = 12
 \end{aligned}$$

Indeed, it aligns with the test from `test/Math.witness.test.ts`:

test/Math.witness.test.ts

```

expect(witness[0]).to.equal(1n);
expect(witness[1]).to.equal(12n); // r
expect(witness[2]).to.equal(1n); // x1
expect(witness[3]).to.equal(3n); // x2
expect(witness[4]).to.equal(4n); // x3
expect(witness[5]).to.equal(12n); // mult
expect(witness[6]).to.equal(12n); // selectMult

```

The initial 1 in the witness is a constant to facilitate the usage of constants inside the circuit. This corresponds to the fact that $w_0 = 1$ is often used to handle constant terms in R1CS.

The Circom also provides a named representation of all witness elements, which is stored in the `.sym` file and looks as follows:

.sym file for $x_1? x_2 \times x_3 : x_2 + x_3$

```

1,1,0,main.r
2,2,0,main.x1
3,3,0,main.x2
4,4,0,main.x3
5,5,0,main.mult
6,6,0,main.selectMult

```

It not only tells us the names of all signals, but also includes information about the optimized signals.

Recall the previous example where we used division and `<--` to store it in the intermediate signal. The generated `.sym` file would look like this:

.sym file for $b \leftarrow x/y, r \leq by$

```

1,1,0,main.r
2,-1,0,main.x
3,2,0,main.y
4,3,0,main.b

```

As we can see, the `-1` was added to the signal `x` to indicate that it is not used in the witness.

Also, according to the documentation of Circom, there are three levels of optimization (ref: <https://docs.circom.io/getting-started/compilation-options>).

In the 2.1.9 version of Circom, the default optimization was `O2`, but it was lowered to `O1` in following versions because `O2` was too aggressive leading to vulnerable circuits.

Remark. In addition, all linear constraints are optimized on the `O2` optimization.

Now, let's examine what the third column in the `.sym` file means.

Consider the following circuit:

```

1 template BinaryCheck() {
2   signal input x1;
3
4   x1 * x1 === x1;
5 }
6

```

```

7  template SelectMult() {
8      signal input x1;
9
10     signal input x2;
11     signal input x3;
12
13     signal mult <== x2 * x3;
14
15     signal output out <== x1 * mult;
16 }
17
18 template Math() {
19     signal output r;
20
21     signal input x1;
22
23     signal input x2;
24     signal input x3;
25
26     component binCheck = BinaryCheck();
27     binCheck.x1 <== x1;
28
29     component selectMult = SelectMult();
30     selectMult.x1 <== x1;
31     selectMult.x2 <== x2;
32     selectMult.x3 <== x3;
33
34     (1 - x1) * (x2 + x3) + selectMult.out ==> r;
35 }

```

We split the circuit into three parts, and the `.sym` file will look like this:

.sym file for $x_1? x_2 \times x_3 : x_2 + x_3$ with templates

```

1,1,2,main.r
2,2,2,main.x1
3,3,2,main.x2
4,4,2,main.x3
5,-1,0,main.binCheck.x1
6,5,1,main.selectMult.out
7,-1,1,main.selectMult.x1
8,-1,1,main.selectMult.x2
9,-1,1,main.selectMult.x3
10,6,1,main.selectMult.mult

```

As we can see, the third column indicates the locality of the signal. Essentially, it tells us which signals are grouped together under the same template.

Another interesting aspect is the order: 0 represents the first component, 1 represents the second component used during computation, and the last component used is the actual 'main'.

Remark. Pay attention to the difference between the modified `Math` circuit's `.sym` file and the original one. Even though we added more signals (i.e., constraints), they were actually optimized by Circom back to the original state.

And the last column is the full name of the constraint, including the path to where it is defined in the code.

11.1.6 Parallel and Custom Keywords

In Circom, there are two special keywords designed to address specific use cases: `custom` and `parallel`.

The `custom` keyword introduces *custom templates* that do not emit R1CS constraints directly. Instead, they delegate logic to `snarkjs` or other libraries at a later stage. Consequently, custom templates **cannot** declare subcomponents or add R1CS constraints within their bodies.

The `custom` keyword is used as follows:

```

1 pragma circom 2.0.6;
2 pragma custom_templates;
3
4 template custom MyCustomGate() {
5     // Custom template's code
6     // No R1CS constraints or subcomponents can be declared here
7     // Logic will be handled by snarkjs as a PLONK custom gates
8 }
```

Remark. At the moment of writing the document the `snarkjs` does not support any custom gates (as stated in their documentation). Also, they can be used only in turbo-PLONK or UltraPlonk schemes. Nevertheless, you can find an example of how they have been used here: <https://github.com/zkFHE/circomlib-fhe/tree/main>.

Meanwhile, the `parallel` keyword (available from Circom 2.0.8 onward) can be applied at either the template or the component instantiation level to parallelize witness generation for independent computations, thereby accelerating large circuits. Parallelism is *only* applied to the C++ witness generator; it does not affect the constraints themselves.

This keyword will be useful in the structures as below:

```

1 template parallel ParallelExample(n) {
2     signal input in[n];
3     signal output out[n];
4
5     // Each iteration is independent, so we can parallelize
6     for (var i = 0; i < n; i++) {
7         out[i] <== in[i] * 2;
8     }
9 }
```

In summary, you should use the `custom` keyword whenever you want to define a template handled **only** by turbo-PLONK or UltraPlonk schemes. You can also find the exact section in the R1CS binary format where custom gates are stored for later processing by the library here: https://github.com/iden3/r1csfile/blob/master/doc/r1cs_bin_format.md#custom-gates-list-section-plonk

The `parallel` keyword is helpful when dealing with large or repetitive computations, as it can speed up witness generation. However, in small circuits or wherever computation is inherently sequential (i.e., where the output of one part is the input to another), `parallel` has no effect.

You can find additional examples of the `parallel` keyword usage here: <https://github.com/zkFHE/circomlib-fhe/tree/main>.

With this, we have covered all the important files generated by Circom.

11.1.7 Generating and Verifying Proofs

Now, it is time to look at proof generation and verification. In this chapter, our main focus will be on the code from `test/Math.circuit.ts`.

To generate a proof, we need to call the `generateProof` method on the circuit object:

```
const proof = await circuit.generateProof(inputs);
```

The actual proof looks like this:

proof.json

```
{
  "proof": {
    "pi_a": [
      "4705801711565477046837119510773988173091957417270766918367441244292047980064",
      "1400811599548904237959319989696481634963162026439383059052135976273120564167",
      "1"
    ],
    "pi_b": [
      [
        "12538508168416900299033726521685163817792614632620657244409429354131980454661",
        "10914283679966848917795247355212516197618338956682374874239005506750384424444"
      ],
      [
        "11504632457518572930719312464170675169899321263873993433191427524966381618623",
        "15524163713890313070296837080299781036987071183397727452907670321368057103914"
      ],
      [
        "1",
        "0"
      ]
    ],
    "pi_c": [
      "260996700533282086084038116247679709285710726946875725263543647585988798998",
      "14278428069254250939292704696175748719031859166075451182707331713513969403299",
      "1"
    ],
    "protocol": "groth16",
    "curve": "bn128"
  },
  "publicSignals": {
    "r": "18"
  }
}
```

Also, at the end of the proof, we have the public signals.

Remark. Usually, public signals are represented by an array of elements, but when using the `hardhat-zkit` plugin, they are typed, and we have actual names for them.

The third element of each program does not participate in any computations; it is needed as additional metadata for the library that implements Groth16 verification.

Remark. When submitting the proof, we have to swap elements inside the arrays of the `b` point, so that the proof can be verified correctly.

These three points π_L, π_R, π_O are used by the verifier to check the equality:

$$e(\pi_L, \pi_R) = e(g_1^\alpha, g_2^\beta) e(\pi_{io}, g_2^\gamma) e(\pi_O, g_2^\delta).$$

Other constants needed for the verifier (for example, points g_1^α or g_2^β) are defined in the following file: `zkit/artifacts/circuits/Math.circom/Math.vkey.json` — see [Figure 11.2](#).

Quick reminder about the structure of points in the proof for BN254 (BN128) curve:

- Each point is either from the regular curve $\mathbb{G}_1 : y^2 = x^3 + b$ over \mathbb{F}_p or from the quadratic extension curve $\mathbb{G}_2 : y'^2 = x'^3 + b'$ over \mathbb{F}_{p^2} . For BN254, the quadratic extension is defined as $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$ with $i^2 + 1$. Curve coefficients are $b = 3 \in \mathbb{F}_p$ and $b' = \frac{3}{9+i} \in \mathbb{F}_{p^2}$.
- Left inputs to the pairing function e are the points on the regular curve \mathbb{G}_1 . They are specified in the form of two field elements $(x, y) \in \mathbb{G}_1$, where $x, y \in \mathbb{F}_p$ are the coordinates.
- Right inputs to the pairing function e are the points over the quadratic extension curve \mathbb{G}_2 . They are specified of the form of four prime field elements $(x_1, y_1, x_2, y_2) \in \mathbb{G}_2$, where the coordinates are $x_1 + iy_1, x_2 + iy_2 \in \mathbb{F}_{p^2}$.
- $e(g_1^\alpha, g_2^\beta)$ is the element from the multiplicative group $\mathbb{F}_{p^{12}}^\times$. Therefore, we need 12 prime field elements to represent it.

Remark (On representing $\mathbb{F}_{p^{12}}$ element). One might wonder: why is the element from $\mathbb{F}_{p^{12}}$ is represented as a pair of two arrays, each consisting of three pairs of prime field elements? The primary reason is that the most convenient way to construct $\mathbb{F}_{p^{12}}$ element is to use the so-called **tower of extensions**: we represent an element from $\mathbb{F}_{p^{12}}$ as a pair of two \mathbb{F}_{p^6} elements, while each \mathbb{F}_{p^6} consists of a triplet of \mathbb{F}_{p^2} elements. For more details, see [Section 3](#)

Thus, we have covered all the information about the internal structure of the Circom files needed for proof generation and verification.

Finally, we verify the proof in the code:

```
expect(await math.verifyProof(proof)).to.be.true
```

This concludes our first journey into learning Circom.

Remark. Here is a set of links that can be used for a deeper dive into the Circom ecosystem:

1. Circom Documentation: <https://docs.circom.io/>
2. Circom Libraries (like circomlib): <https://github.com/iden3/circomlib>

vkey.json

```

{
  "protocol": "groth16",
  "curve": "bn128",
  "nPublic": 1,
  "vk_alpha_1": [
    "20491192805390485299153009773594534940189261866228447918068658471970481763042",
    "9383485363053290200918347156157836566562967994039712273449902621266178545958",
    "1"
  ],
  "vk_beta_2": [
    [
      "6375614351688725206403948262868962793625744043794305715222011528459656738731",
      "4252822878758300859123897981450591353533073413197771768651442665752259397132"
    ],
    [
      "10505242626370262277552901082094356697409835680220590971873171140371331206856",
      "21847035105528745403288232691147584728191162732299865338377159692350059136679"
    ],
    [ "1", "0" ]
  ],
  "vk_gamma_2": [
    [
      "10857046999023057135944570762232829481370756359578518086990519993285655852781",
      "11559732032986387107991004021392285783925812861821192530917403151452391805634"
    ],
    [
      "8495653923123431417604973247489272438418190587263600148770280649306958101930",
      "4082367875863433681332203403145435568316851327593401208105741076214120093531"
    ],
    [ "1", "0" ]
  ],
  "vk_delta_2": [
    [
      "10857046999023057135944570762232829481370756359578518086990519993285655852781",
      "11559732032986387107991004021392285783925812861821192530917403151452391805634"
    ],
    [
      "8495653923123431417604973247489272438418190587263600148770280649306958101930",
      "4082367875863433681332203403145435568316851327593401208105741076214120093531"
    ],
    [ "1", "0" ]
  ],
  "vk_alphabeta_12": [
    [
      "2029413683389138792403550203267699914886160938906632433982220835551125967885",
      "21072700047562757817161031222997517981543347628379360635925549008442030252106"
    ],
    [
      "5940354580057074848093997050200682056184807770593307860589430076672439820312",
      "12156638873931618554171829126792193045421052652279363021382169897324752428276"
    ],
    [
      "7898200236362823042373859371574133993780991612861777490112507062703164551277",
      "7074218545237549455313236346927434013100842096812539264420499035217050630853"
    ],
    [
      "7077479683546002997211712695946002074877511277312570035766170199895071832130",
      "10093483419865920389913245021038182291233451549023025229112148274109565435465"
    ],
    [
      "4595479056700221319381530156280926371456704509942304414423590385166031118820",
      "19831328484489333784475432780421641293929726139240675179672856274388269393268"
    ],
    [
      "11934129596455521040620786944827826205713621633706285934057045369193958244500",
      "8037395052364110730298837004334506829870972346962140206007064471173334027475"
    ],
    [
      "4162541565828872643496914921393902054824387648641933177665940781539334781623",
      "4678293780284819015763290392952715769540194300841323348855962545628746384938",
      "1"
    ],
    [
      "7846408072049176620553358542204120795817938985459251067840222635524693287955",
      "17494754572705064819681843056808269434754047134538681907566256240907807975850",
      "1"
    ],
    [
      "1"
    ]
  ],
  "IC": [
    [
      "4162541565828872643496914921393902054824387648641933177665940781539334781623",
      "4678293780284819015763290392952715769540194300841323348855962545628746384938",
      "1"
    ],
    [
      "7846408072049176620553358542204120795817938985459251067840222635524693287955",
      "17494754572705064819681843056808269434754047134538681907566256240907807975850",
      "1"
    ],
    [
      "1"
    ]
  ]
}

```

Figure 11.2: Verification key for the proof stored in the generated vkey.json file.