



# STARK by Hand

This document offers a concrete example of the RISC Zero Proof System; you can find a PDF version on [www.RISCZero.com](http://www.RISCZero.com) or peek behind the formulas with the [Google Sheet Version](#).

For questions, corrections, conversation, and collaboration, find us on [Twitter](#) or [Discord](#).

When any code executes in the RISC Zero virtual machine, each step of that execution is recorded in an **Execution Trace**. We show a simplified example, computing 4 steps of a Fibonacci sequence modulo 97, using two user-specified inputs. In this sheet, we introduce the columns of a RISC Zero Execution Trace. In the following sheets, we will demonstrate a concrete example of how RISC Zero proves the validity of an Execution Trace without revealing any knowledge.

In this example, our trace consists of 6 columns. Each of the first three columns is a record of the internal state of a register at each clock cycle from initialization until termination. We call these **Data Columns**. The next three columns are **Control Columns**, which we use to mark initialization and termination points.

In the full RISC Zero protocol, the **Data Columns** hold the state of the RISC-V processor, including ISA registers, the program counter, and various microarchitecture details such as instruction decoding data, ALU registers, etc., while the **Control Columns** handle system initialization and shutdown, the initial program code to load into memory before execution, and other control signals that don't depend on the programs execution.

Input 1	24
Input 2	30

Asserted Output	28
-----------------	----

Modulo	97
--------	----

	Clock Cycle	Data Column 1	Data Column 2	Data Column 3	Control Column - Initialization	Control Column - Transition	Control Column - Termination
Execution Trace - Initialization	0	24	30	54	1	0	0
Execution Trace - Transition	1	30	54	84	0	1	0
Execution Trace - Transition	2	54	84	41	0	1	0
Execution Trace - Termination	3	84	41	28	0	1	1



## Checking the Trace

In this sheet, we introduce a number of rule-checking cells in order to demonstrate the validity of the Execution In this example, we show six rules. In the full RISC-V implementation, we check over 100 rules in order to validate the execution trace.

Each rule check is written as the product of two terms, modulo 97. The first term equals zero when the rule holds. The second term equals zero when we don't want to enforce the rule.

Each rule checking column can be expressed as a multi-input, single-output polynomial, where the inputs are some combination of entries in the trace; we call these **Rule-Checking Polynomials**.

Input 1	24
Input 2	30

Asserted Output	28
-----------------	----

Modulo	97
--------	----

	Clock Cycle	Data Column 1	Data Column 2	Data Column 3	Control Column - Initialization	Control Column - Transition	Control Column - Termination	Does Fibonacci relation hold?	Does Initialization for Data Column 1 match User Input 1?	Does Initialization for Data Column 2 match User Input 2?	Does Termination value for Data Column 3 match Output?	Does entry i from Input Column 1 match entry i-1 from Input Column 2?	Does entry i from Input Column 2 match entry i-1 from Output Column?
Execution Trace - Initialization	0	24	30	54	1	0	0	0	0	0	0	0	0
Execution Trace - Transition	1	30	54	84	0	1	0	0	0	0	0	0	0
Execution Trace - Transition	2	54	84	41	0	1	0	0	0	0	0	0	0
Execution Trace - Termination	3	84	41	28	0	1	1	0	0	0	0	0	0





# Interpolating Trace Polynomials

Let's remove the rule-checking columns for a minute and turn our attention toward encoding our Trace data in terms of polynomials. Just as any two points allow you to draw a line, any 8 points allow you to determine a degree 7 polynomial. The standard techniques for constructing a polynomial through a given set of points are Lagrange Interpolation and Finite Fourier Transforms (FFTs). In our context, we use NTTs (Number Theoretic Transforms), which are essentially just a finite field equivalent of an FFT. Since we're working modulo 97, our polynomials have values and coefficients in the finite field  $F_{97}$ .

The polynomial encoding technique we use in RISC Zero is called Reed-Solomon encoding; RS codes are ubiquitous in the world of digital signal processing as a method of providing redundancy for error checking and error correction purposes.

Reed-Solomon Codes are built using points of the form  $(a^0, x_0)$ ,  $(a^1, x_1)$ , ... In this example, we use powers of 28 as the inputs (and the entries of the trace as the outputs).

We use this technique to write a **Trace Polynomial** for each column of the trace. Running an iNTT on the 8 entries from Data Column 1, we use  $iNTT(\text{column}, \text{modulus})$  to generate a polynomial whose evaluations agree with the trace data. Then, we evaluate this polynomial over an **Expanded Domain** to construct the Reed Solomon encoding of the column.

In Python using sympy, `intt([24, 30, 54, 84, 78, 15, 29, 50], prime=97)` returns `[94, 68, 41, 69, 25, 72, 85, 55]`.

The 8 entries of this iNTT input array are shown in **boldface** in Data Column 1 below; with each entry corresponding to a row of the Padded Trace.

We use the entries of the output array as the coefficients of the associated Trace Polynomial. In this case,  $d_1(x) = 94 + 68x + 41x^2 + 69x^3 + 25x^4 + 72x^5 + 85x^6 + 55x^7 \pmod{97}$ .

The key feature of  $d_1(x)$  is that for  $z=5^0, 5^12, 5^24, \dots$  the evaluations  $d_1(z)$  agree with the values in Data Column 1:  $d_1(5^0) = 24$ ,  $d_1(5^12) = 30$ ,  $d_1(5^24) = 54$ , etc.

We proceed similarly on each data column and each control column, generating one Trace Polynomial for each column of our Trace.

A quick note about  $F_{97}$ : every element of this field can be written as a power of 5. In other words, the elements of  $F_{97}$  are  $0, 5^0, 5^1, \dots$  and  $5^95$ . Written in this form, we can view our Reed-Solomon inputs as every third power of 5:  $5^0, 5^3, 5^6$ , etc.

For a brief introduction to finite fields as they relate to the RISC Zero proof system, see [here](#).

Input 1	24
Input 2	30
Modulo	97

Asserted Output	28
-----------------	----

	Reed Solomon Input (Exponent Notation in terms of 28)	Reed Solomon Input (Exponent Notation in terms of 5)	Reed Solomon Input (Simplified)	Clock Cycle	Data Column 1	Data Column 2	Data Column 3	Control Column - Initialization	Control Column - Transition	Control Column - Termination
<b>Execution Trace - Initialization</b>	<b>28<sup>0</sup> mod 97</b>	<b>5<sup>0</sup> mod 97</b>	<b>1</b>	<b>0</b>	<b>24</b>	<b>30</b>	54	1	0	0
RS Redundancy Row	28 <sup>1</sup> mod 97	5 <sup>3</sup> mod 97	28		27	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row	28 <sup>2</sup> mod 97	5 <sup>6</sup> mod 97	8		74	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row	28 <sup>3</sup> mod 97	5 <sup>9</sup> mod 97	30		77	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
<b>Execution Trace - Transition</b>			<b>64</b>	<b>1</b>	<b>30</b>	54	84	0	1	0
RS Redundancy Row			46		37	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			27		62	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			77		3	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
<b>Execution Trace - Termination</b>			<b>22</b>	<b>2</b>	<b>54</b>	84	41	0	1	0
RS Redundancy Row			34		42	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			79		96	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			78		69	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
<b>Execution Trace - Termination</b>			<b>50</b>	<b>3</b>	<b>84</b>	41	<b>28</b>	0	0	1
RS Redundancy Row			42		36	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			12		26	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			45		37	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
<b>Random Padding</b>			<b>96</b>	<b>4</b>	<b>78</b>	95	77	0	0	0
RS Redundancy Row			69		71	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			89		24	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			67		70	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
<b>Random Padding</b>			<b>33</b>	<b>5</b>	<b>15</b>	52	7	0	0	0
RS Redundancy Row			51		31	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			70		38	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			20		71	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
<b>Random Padding</b>			<b>75</b>	<b>6</b>	<b>29</b>	82	12	0	0	0
RS Redundancy Row			63		40	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			18		54	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			19		19	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
<b>Random Padding</b>			<b>47</b>	<b>7</b>	<b>50</b>	12	26	0	0	0
RS Redundancy Row			55		80	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row			85		87	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row	28 <sup>31</sup> mod 97	5 <sup>93</sup> mod 97	52		18	For each column, compute iNTT(Trace Column, modulus) to fill in this data				
RS Redundancy Row	28 <sup>32</sup> mod 97	5 <sup>96</sup> mod 97	1		24	For each column, compute iNTT(Trace Column, modulus) to fill in this data				

97  
5 order 96  
we have a trace of length 8  
5<sup>12</sup>  
5<sup>3</sup>  
x0, x1, x2, x3, x4, ..., x7  
(g<sup>0</sup>, x0), (g<sup>1</sup>, x1), ... (g<sup>7</sup>, x7)



## Interpolating Trace Polynomials (cont.)

This sheet shows our 6 trace polynomials, each evaluated at 32 points.

intt(TraceColumn, prime=97) returns the coefficients of the trace polynomial.

Column	Python Code (from sympy import intt)	Code Output (Coefficients of Trace Polynomials)							
d1	d1 = intt([24, 30, 54, 84, 78, 15, 29, 50], prime=97)	94	68	41	69	25	72	85	55
d2	d2 = intt([30, 54, 84, 41, 2, 77, 21, 36], prime=97)	31	31	0	87	76	66	6	24
d3	d3 = intt([54, 84, 41, 28, 71, 17, 92, 33], prime=97)	4	14	83	44	12	44	12	35
c1	c1 = intt([1, 0, 0, 0, 0, 0, 0], prime=97)	85	85	85	85	85	85	85	85
c2	c2 = intt([0, 1, 1, 1, 0, 0, 0], prime=97)	61	80	12	37	12	60	12	17
c3	c3 = intt([0, 0, 0, 1, 0, 0, 0], prime=97)	85	89	27	18	12	8	70	79

Trace Polynomials
d1(x)=94+68x+41x <sup>2</sup> +69x <sup>3</sup> +25x <sup>4</sup> +72x <sup>5</sup> +85x <sup>6</sup> +55x <sup>7</sup>
d2(x)=31+31x+0x <sup>2</sup> +87x <sup>3</sup> +76x <sup>4</sup> +66x <sup>5</sup> +6x <sup>6</sup> +24x <sup>7</sup>
d3(x)=4+14x+83x <sup>2</sup> +44x <sup>3</sup> +12x <sup>4</sup> +44x <sup>5</sup> +12x <sup>6</sup> +35x <sup>7</sup>
c1(x)=85+85x+85x <sup>2</sup> +85x <sup>3</sup> +85x <sup>4</sup> +85x <sup>5</sup> +85x <sup>6</sup> +85x <sup>7</sup>
c2(x)=61+80x+12x <sup>2</sup> +37x <sup>3</sup> +12x <sup>4</sup> +60x <sup>5</sup> +12x <sup>6</sup> +17x <sup>7</sup>
c3(x)=85+89x+27x <sup>2</sup> +18x <sup>3</sup> +12x <sup>4</sup> +8x <sup>5</sup> +70x <sup>6</sup> +79x <sup>7</sup>

Input 1	24
Input 2	30

Asserted Output	28
-----------------	----

Modulo	97
--------	----

	Reed Solomon Input (Exponent Notation in terms of 28)	Reed Solomon Input (Exponent Notation in terms of 5)	Reed Solomon Input (Simplified)	Clock Cycle	Data Column 1	Data Column 2	Data Column 3	Control Column - Initialization	Control Column - Transition	Control Column - Termination
<b>Execution Trace - Initialization</b>	28 <sup>0</sup> mod 97	5 <sup>0</sup> mod 97	1	0	24	30	54	1	0	0
RS Redundancy Row	28 <sup>1</sup> mod 97	5 <sup>3</sup> mod 97	28		27	33	43	23	35	26
RS Redundancy Row	28 <sup>2</sup> mod 97	5 <sup>6</sup> mod 97	8		74	14	27	45	31	13
RS Redundancy Row	28 <sup>3</sup> mod 97	5 <sup>9</sup> mod 97	30		77	31	88	53	63	86
<b>Execution Trace - Transition</b>			64	1	30	54	84	0	1	0
RS Redundancy Row			46		37	76	67	72	32	46
RS Redundancy Row			27		62	85	34	83	63	87
RS Redundancy Row			77		3	84	63	70	30	80
<b>Execution Trace - Transition</b>			22	2	54	84	41	0	1	0
RS Redundancy Row			34		42	14	11	10	58	60
RS Redundancy Row			79		96	18	86	60	59	28
RS Redundancy Row			78		69	3	29	25	20	91
<b>Execution Trace - Termination</b>			50	3	84	41	28	0	1	1
RS Redundancy Row			42		36	15	54	53	8	23
RS Redundancy Row			12		26	16	31	11	91	45
RS Redundancy Row			45		37	77	1	68	51	53
<b>Random Padding</b>			96	4	78	2	71	0	0	0
RS Redundancy Row			69		71	90	82	2	38	72
RS Redundancy Row			89		24	15	14	62	57	83
RS Redundancy Row			67		70	66	75	13	66	70
<b>Random Padding</b>			33	5	15	77	17	0	0	0
RS Redundancy Row			51		31	8	76	26	65	10
RS Redundancy Row			70		38	20	67	13	36	60
RS Redundancy Row			20		71	19	14	86	9	25
<b>Random Padding</b>			75	6	29	21	92	0	0	0
RS Redundancy Row			63		40	43	42	46	81	53
RS Redundancy Row			18		54	72	72	87	86	11
RS Redundancy Row			19		19	92	90	80	70	68
<b>Random Padding</b>			47	7	50	36	33	0	0	0
RS Redundancy Row			55		80	66	45	60	74	2
RS Redundancy Row			85		87	8	89	28	65	62
RS Redundancy Row	28 <sup>31</sup> mod 97		52		18	70	60	91	82	13
RS Redundancy Row	28 <sup>32</sup> mod 97		1							



# Committing Trace Polynomials

The next step is for the Prover to commit the **Trace Polynomials** into a **Merkle Tree**. In order to maintain a Zero-Knowledge protocol, the Prover evaluates each Trace Polynomial over a "shifted evaluation domain." Specifically, we evaluate each  $d_i(x)$  at  $x=5, 5^4, 5^7, \dots, 5^{93}$ .

Note that because of our shifted evaluation domain, the yellow and blue cells in Data Columns 1, 2, and 3 no longer match the Inputs and Asserted Outputs. In fact, this shift in the evaluation domain disguises *all* the Trace Data. We only reveal information about the disguised trace, and the random padding we appended is sufficient to prevent an attacker from deducing any connection between the disguised trace and the actual trace.

Trace Polynomials
$d1(x)=94+68x+41x^2+69x^3=25x^4+72x^5+85x^6+55x^7$
$d2(x)=31+31x+0x^2+87x^3=76x^4+66x^5+6x^6+24x^7$
$d3(x)=4+14x+83x^2+44x^3=12x^4+44x^5+12x^6+35x^7$
$c1(x)=85+85x+85x^2+85x^3=85x^4+85x^5+85x^6+85x^7$
$c2(x)=61+80x+12x^2+37x^3=12x^4+60x^5+12x^6+17x^7$
$c3(x)=85+89x+27x^2+18x^3=12x^4+8x^5+70x^6+79x^7$

Input 1	24
Input 2	30

Asserted Output	28
-----------------	----

Modulo	97
--------	----

	Reed Solomon Input (Exponent Notation in terms of 28)	Reed Solomon Input (Exponent Notation in terms of 5)	Reed Solomon Input (Simplified)	Clock Cycle	Data Column 1	Data Column 2	Data Column 3	Control Column - Initialization	Control Column - Transition	Control Column - Termination
<b>Disguised Execution Trace</b>	$5^{*28^0} \text{ mod } 97$	$5^1 \text{ mod } 97$	5	0	31	39	12	82	81	2
Disguised RS Redundancy Row	$5^{*28^1} \text{ mod } 97$	$5^4 \text{ mod } 97$	43		15	36	11	18	72	32
Disguised RS Redundancy Row	$5^{*28^2} \text{ mod } 97$	$5^7 \text{ mod } 97$	40		96	65	6	32	73	65
Disguised RS Redundancy Row	$5^{*28^3} \text{ mod } 97$	$5^{10} \text{ mod } 97$	53		79	41	88	68	10	22
<b>Disguised Execution Trace</b>			29	1	69	35	49	81	64	32
Disguised RS Redundancy Row			36		31	85	50	41	58	16
Disguised RS Redundancy Row			38		16	41	69	18	40	38
Disguised RS Redundancy Row			94		71	24	89	86	56	50
<b>Disguised Execution Trace</b>			13	2	35	77	46	92	16	47
Disguised RS Redundancy Row			73		10	40	9	59	83	24
Disguised RS Redundancy Row			7		53	54	58	14	20	67
Disguised RS Redundancy Row			2		28	7	95	44	92	35
<b>Disguised Execution Trace</b>			56	3	67	81	80	43	61	82
Disguised RS Redundancy Row			16		26	2	73	31	21	18
Disguised RS Redundancy Row			60		0	54	76	8	64	32
Disguised RS Redundancy Row			31		45	36	80	92	4	68
<b>Disguised   Random Padding</b>			92	4	91	59	35	10	22	81
Disguised RS Redundancy Row			54		94	39	57	71	34	41
Disguised RS Redundancy Row			57		18	82	19	50	40	18
Disguised RS Redundancy Row			44		80	36	44	89	28	86
<b>Disguised   Random Padding</b>			68	5	54	12	61	2	48	92
Disguised RS Redundancy Row			61		39	46	78	32	64	59
Disguised RS Redundancy Row			59		16	16	16	65	72	14
Disguised RS Redundancy Row			3		19	49	95	22	31	44
<b>Disguised   Random Padding</b>			84	6	57	23	47	32	55	43
Disguised RS Redundancy Row			24		74	89	18	16	37	31
Disguised RS Redundancy Row			90		40	96	42	38	26	8
Disguised RS Redundancy Row			95		43	54	72	50	9	92
<b>Disguised   Random Padding</b>			41	7	57	19	90	47	44	10
Disguised RS Redundancy Row			81		75	8	27	24	22	71
Disguised RS Redundancy Row			37		28	34	37	67	56	50
Disguised RS Redundancy Row	$5^{*28^{31}} \text{ mod } 97$		66		96	1	51	35	64	89
Disguised RS Redundancy Row	$28^{*32} \text{ mod } 97$		5							















# FRI Folding

Given a vector commitment, the FRI protocol proves that the commitment corresponds to evaluations of a low-degree polynomial. In this example, we use FRI to prove that the "FRI Polynomial" commitment (from the previous page) has degree at most 7.

The "FRI blow-up factor" here is 4, since the commitment for the FRI polynomial has 32 entries and a degree 7 polynomial has 8 coefficients. This blow-up factor is a consequence of the choice of "rate" for the Reed-Solomon expansion used earlier. A FRI blow-up factor of 4 corresponds to an RS code of rate 1/4.

FRI consists of a commit phase and a query phase. The commit phase consists of  $r$  rounds. In each round, the prover "folds" the previous commitment into a smaller commitment (both in terms of commitment size and polynomial degree).

Here, we show 3 rounds using a folding factor of 2: in each round, the Prover commits to a vector whose length is half that of the previous commitment.

The folding at each round is accomplished by first splitting the coefficients into an even part and an odd part and then mixing the two parts together using verifier-supplied randomness.

This page shows the coefficient form for the polynomials at each round; the commitments at each round are shown on the following page.

<b>f<sub>0</sub></b>	<b>19 + 56x + 34x<sup>2</sup> + 48x<sup>3</sup> + 43x<sup>4</sup> + 37x<sup>5</sup> + 10x<sup>6</sup> + 0x<sup>7</sup></b>
Even Part of f <sub>0</sub>	19 + 34x + 43x <sup>2</sup> + 10x <sup>3</sup>
Odd Part of f <sub>0</sub>	56 + 48x + 37x <sup>2</sup> + 0x <sup>3</sup>

Round 1 Randomness: 12

$$f_1 = [\text{Even Part of } f_0] + 12 \cdot [\text{Odd Part of } f_0]$$

<b>f<sub>1</sub></b>	<b>12 + 28x + 2x<sup>2</sup> + 10x<sup>3</sup></b>
Even Part of f <sub>1</sub>	12 + 2x
Odd Part of f <sub>1</sub>	28 + 10x

Round 2 Randomness: 32

$$f_2 = [\text{Even Part of } f_1] + 32 \cdot [\text{Odd Part of } f_1]$$

<b>f<sub>2</sub></b>	<b>35 + 31x</b>
Even Part of f <sub>2</sub>	35 + x
Odd Part of f <sub>2</sub>	31 + x

Round 3 Randomness: 64

$$f_3 = [\text{Even Part of } f_2] + 64 \cdot [\text{Odd Part of } f_2]$$

<b>f<sub>3</sub></b>	<b>79</b>
----------------------	-----------



# FRI Queries

The Prover has committed to evaluations of  $f_0$  at powers of 28, evaluations of  $f_1$  at powers of  $28^2$ , evaluations of  $f_2$  at powers of  $28^4$ , and evaluations of  $f_3$  at powers of  $28^8$ .

Since 28 has multiplicative order 32, this corresponds to evaluation domains of 32 elements, 16 elements, 8 elements, and 4 elements.

After these commitments are made, the Verifier makes a number of *queries*. The queries serve as a random challenge, testing the legitimacy of the Prover's commitments. Loosely speaking, with a blow-up factor of 4, a single query will catch a cheating Prover 3/4 of the time. In other words, a single query provides 2 bits of security. The RISC Zero zkVM uses 50 queries and a blow-up factor of 4, which amounts to 100 bits of security.

Note that the analysis above is a substantial simplification of the full security analysis; the precise security level is not exactly 100 bits. For a more thorough security analysis, see our [ZKP Whitepaper](#).

The key point about FRI folding is that it can be checked *locally*. For a single query, the Prover provides 2 evaluations from  $f_0$ ,  $f_1$ , and  $f_2$ , and a single evaluation from  $f_3$ .

In particular, if the Verifier requests a query for  $g$ , the Prover sends evaluations for:  $f_0(\pm g)$ ,  $f_1(\pm g^2)$ ,  $f_2(\pm g^4)$ ,  $f_3(g^8)$

The Verifier can check the evaluations are consistent from round-to-round. For example,  $f_1(g^2)$  can be expressed in terms of  $f_0(g)$ ,  $f_0(-g)$ , and the randomness for that round.

For details, we refer readers to the markdown version of this explainer: <https://www.risczero.com/docs/explainers/proof-system/stark-by-hand>

	$f_0$	$f_1$	$f_2$	$f_3$
<b>Evaluation Domain (Powers of 28)</b>	<b><math>19 + 56x + 34x^2 + 48x^3 + 43x^4 + 37x^5 + 10x^6 + 0x^7</math></b>	<b><math>12 + 28x + 2x^2 + 10x^3</math></b>	<b><math>35 + 31x</math></b>	<b>79</b>
1	53	52	66	79
28	69			
8	63	52		
30	30			
64	46	20	79	
46	13			
27	60	12		
77	50			
22	38	18	38	79
34	3			
79	95	36		
78	23			
50	75	68	33	
42	39			
12	62	68		
45	19			
96	62	73	4	79
69	58			
89	41	34		
67	67			
33	89	92	88	
51	41			
70	50	18		
20	24			
75	95	2	32	79
63	90			
18	72	23		
19	20			
47	82	62	37	
55	33			
85	0	47		
52	16			