Tracing Transactions **Across** Cryptocurrency Ledgers

1

换币、混币

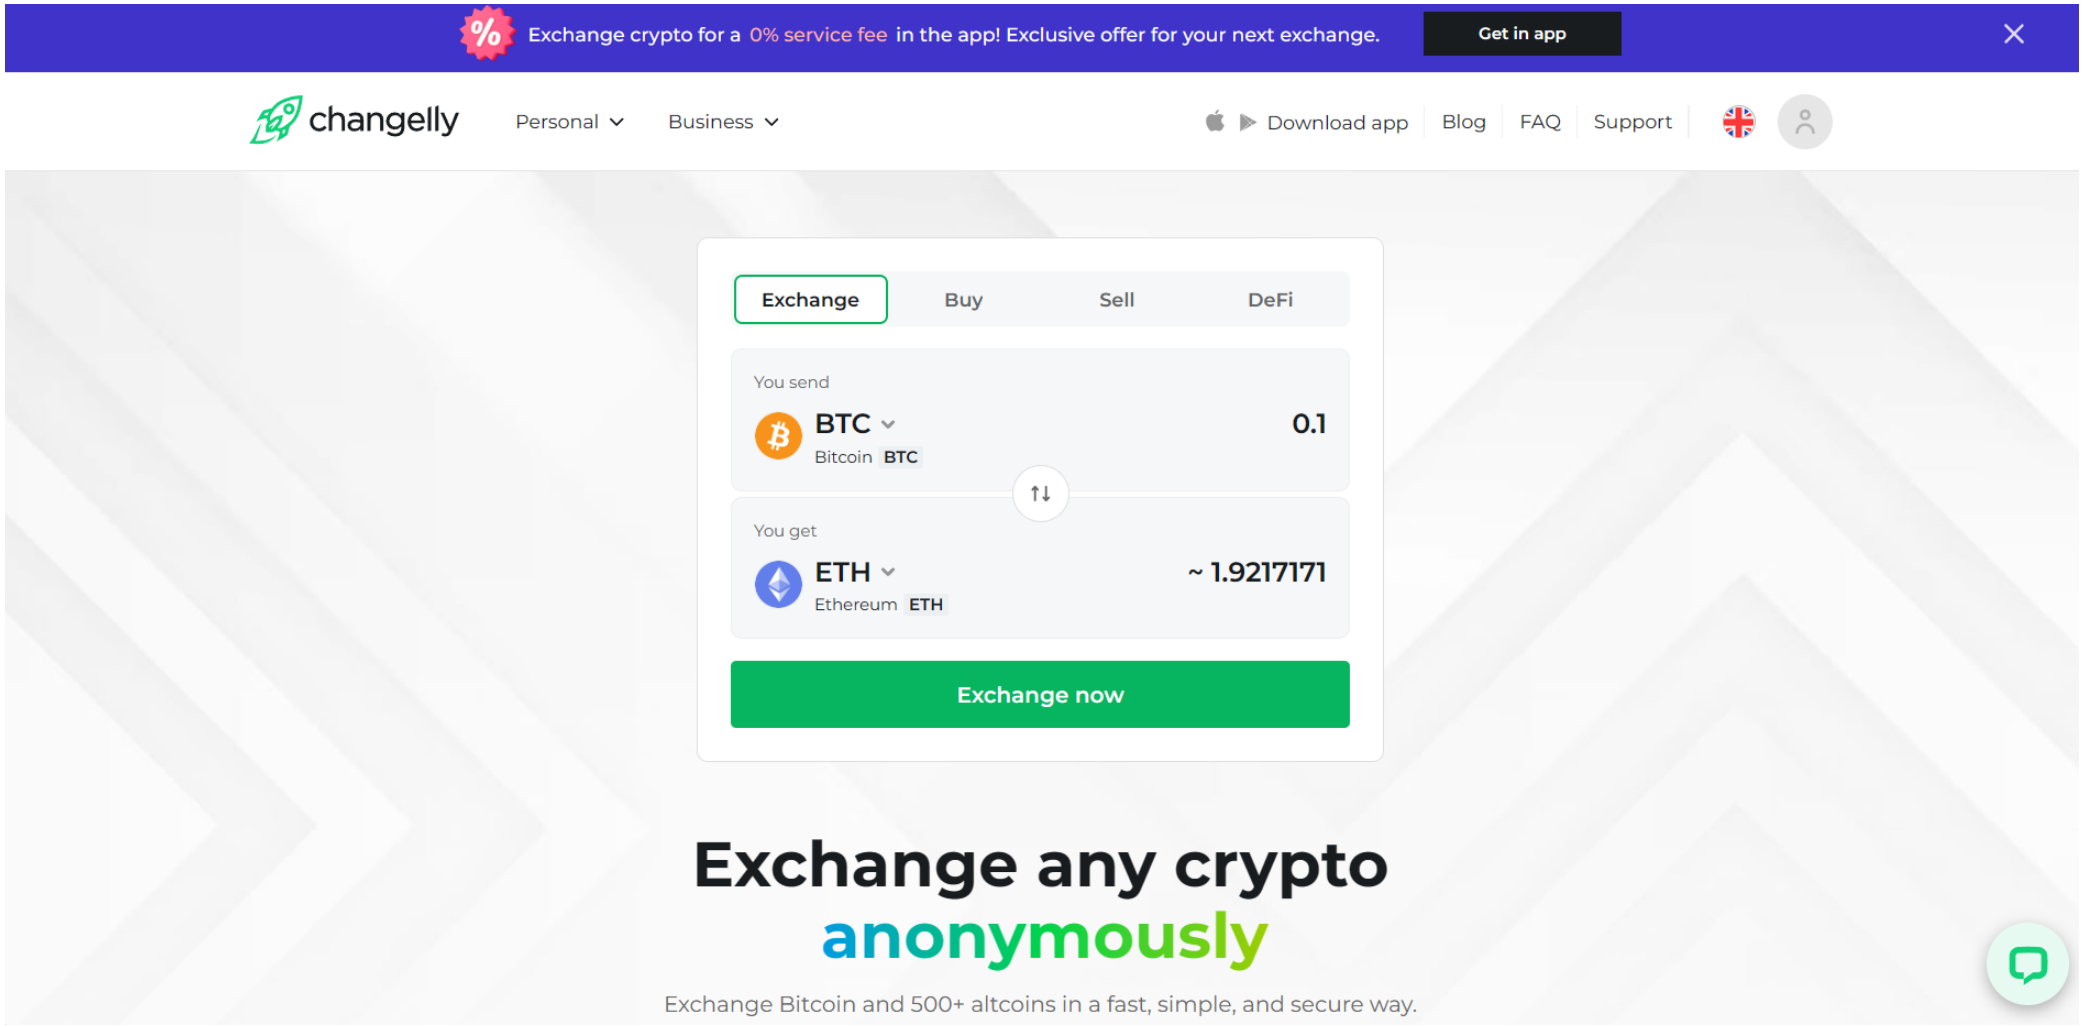
Exchange crypto for a 0% service fee in the app! Exclusive offer for your next exchange.
Get in app
changelly
Personal
Business
Download app
Blog
FAQ
Support
Exchange
Buy
Sell
DeFi
You send
BTC
Bitcoin BTC
0.1
You get
ETH
Ethereum ETH
~ 1.9217171
Exchange now
Exchange any crypto
anonymously
Exchange Bitcoin and 500+ altcoins in a fast, simple, and secure way.

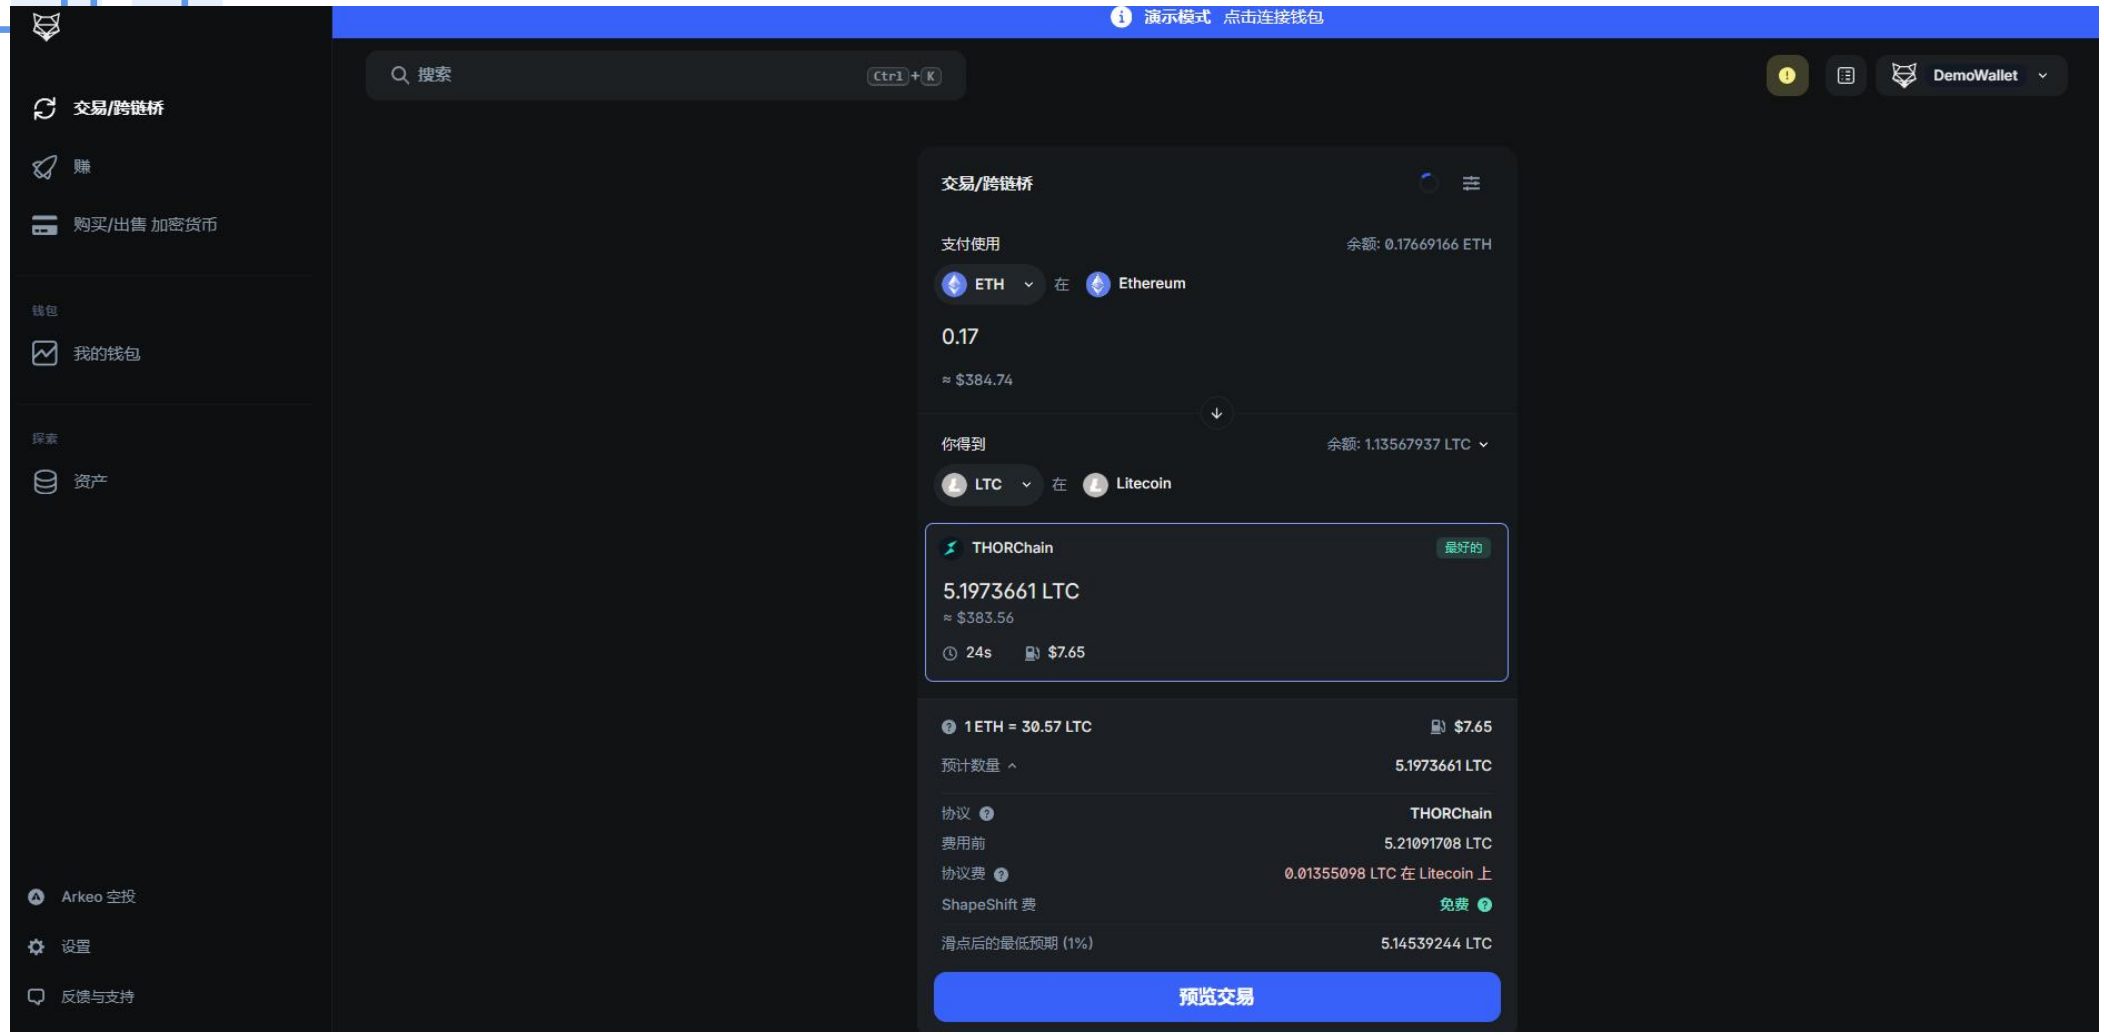
演示模式 点击连接钱包
搜索
Ctrl + K
DemoWallet
交易/跨链桥
赚
购买/出售 加密货币
钱包
我的钱包
探索
资产
Arkeo 空投
设置
反馈与支持
交易/跨链桥
支付使用
余额: 0.17669166 ETH
ETH
在
Ethereum
0.17
≈ $384.74
你得到
余额: 1.13567937 LTC
LTC
在
Litecoin
THORChain
最好的
5.1973661 LTC
≈ $383.56
24s
$7.65
1 ETH = 30.57 LTC
$7.65
预计数量
5.1973661 LTC
协议
THORChain
费用前
5.21091708 LTC
协议费
0.01355098 LTC 在 Litecoin 上
ShapeShift 费
免费
滑点后的最低预期 (1%)
5.14539244 LTC
预览交易


确认详情
发送
0.17 ETH
≈ $384.74
预计数量
5.1973661 LTC
≈ $383.57
利率
1 ETH = 30.57 LTC
@THORChain
矿工费
0.0033807162806 ETH ≃ $7.65
THORChain 交易
这个兑换可能需要几分钟才能完成。
确认并交易

What is a CoinJoin?

In essence, a CoinJoin involves the combination of inputs by multiple users into a single transaction.

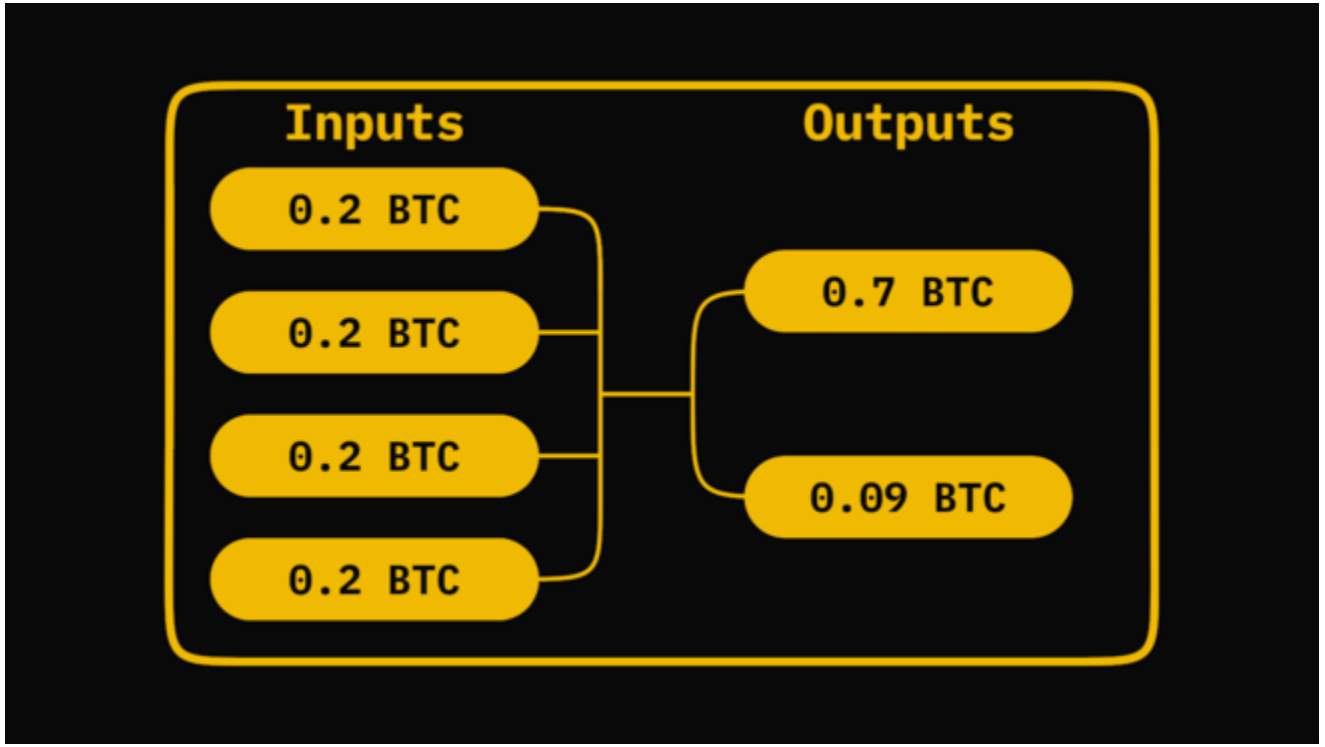

How does a CoinJoin work?

The idea is that multiple parties will coordinate to create a transaction, each providing inputs and desired outputs. As all of the inputs are combined, it becomes impossible to say with certainty which output belongs to which user.

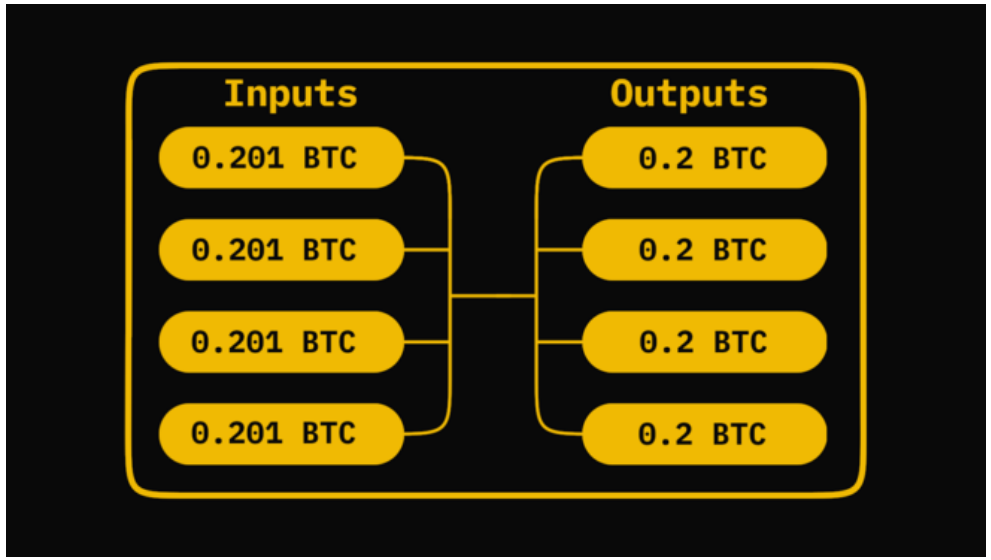

What For?

Anonymity:

- money laundering
- phishing/typosquatting scam
- other illicit purposes

Other Purpose:

- Trading bots

2

基本模型

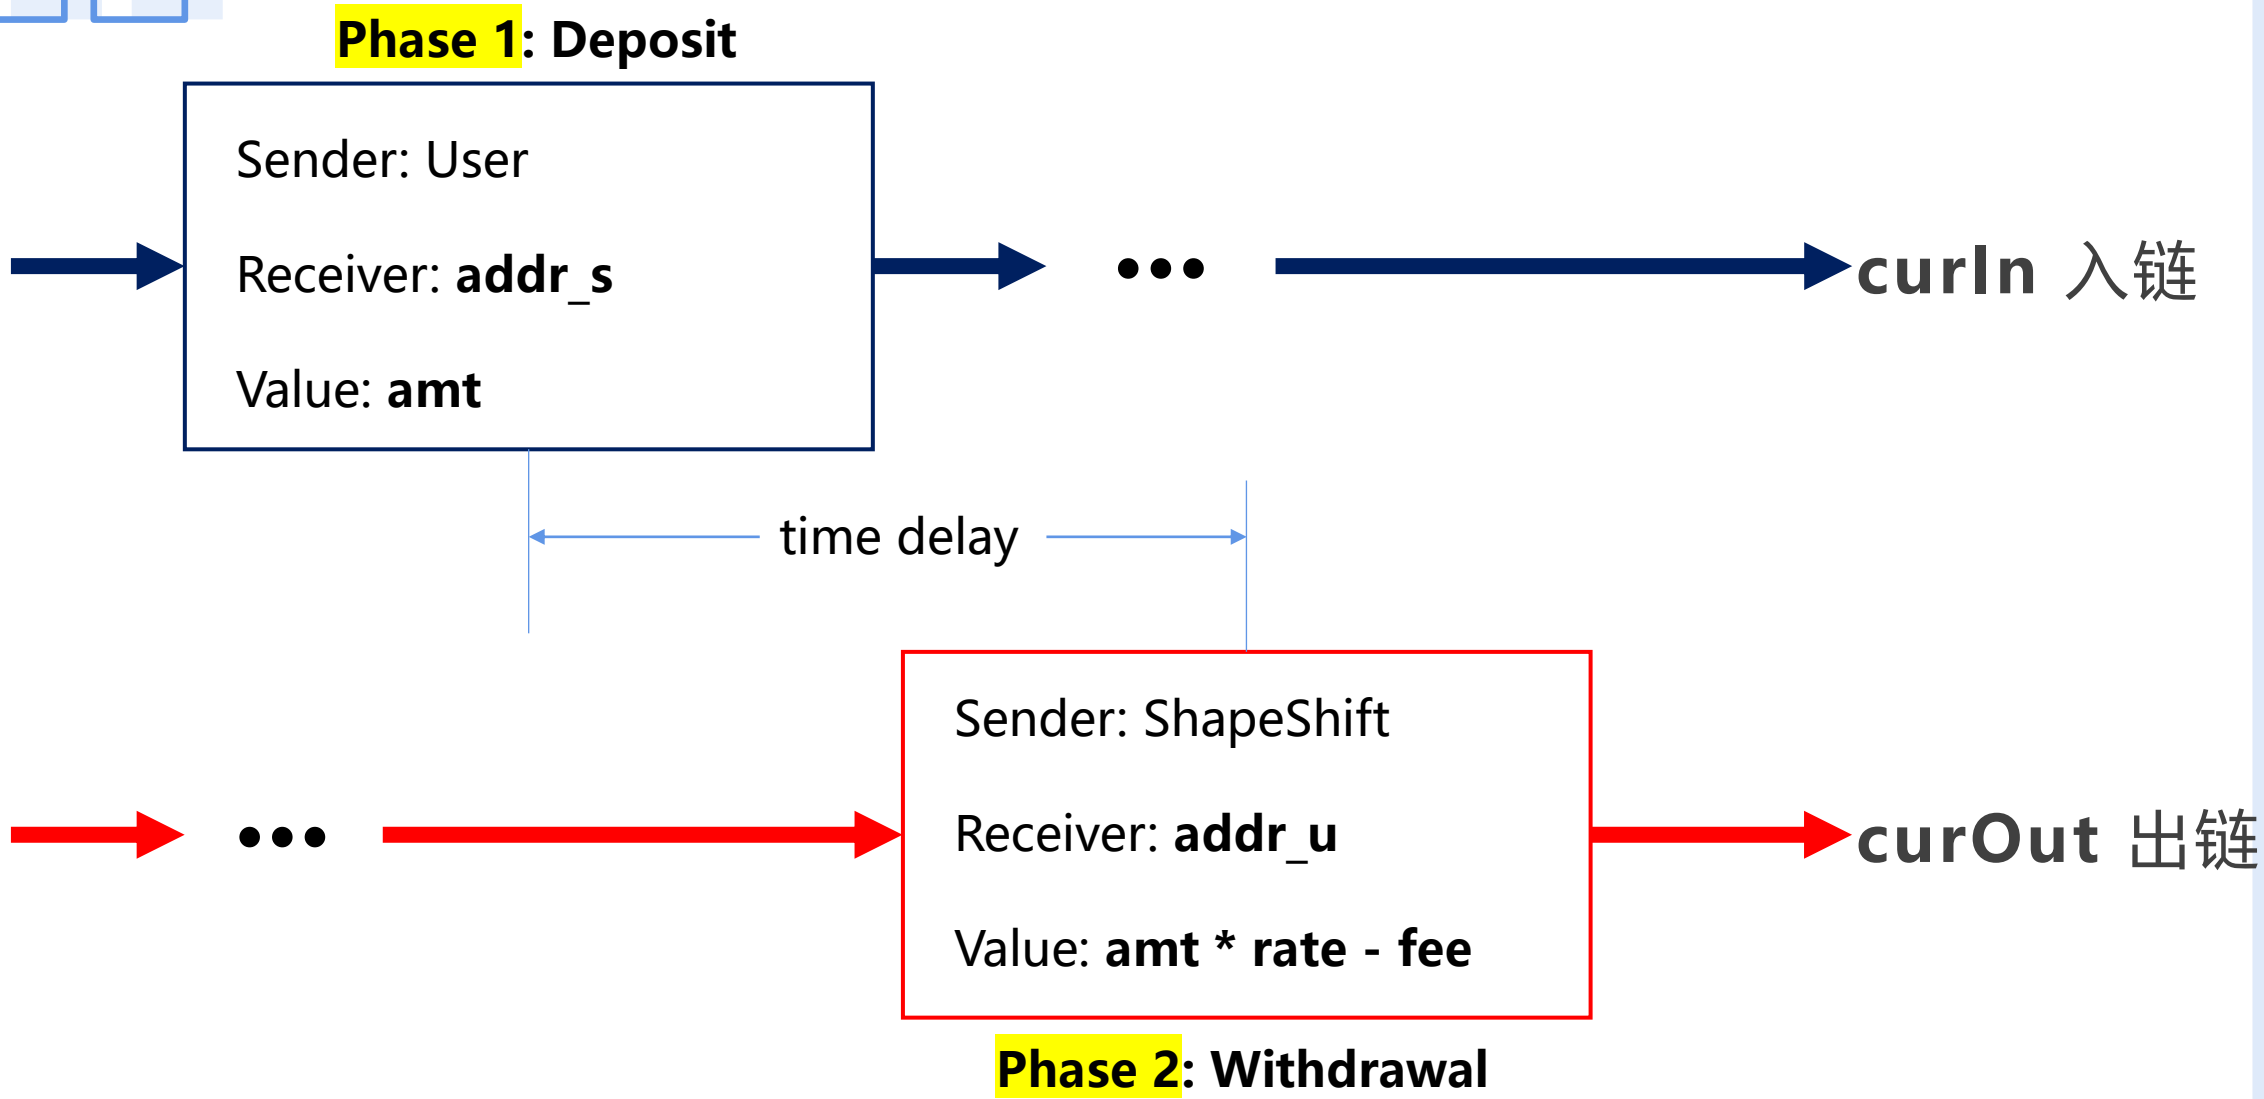
Phase 1: Deposit
Sender: User
Receiver: addr_s
Value: amt
curIn 入链
time delay
Sender: ShapeShift
Receiver: addr_u
Value: amt * rate - fee
curOut 出链
Phase 2: Withdrawal

3

数据获取

数据源

Scrape from exchanger’s public API:

(1) the current trading rate between any pair of cryptocurrencies

(2) a list of up to 50 of the most recent transactions that have taken place (across all users)

(3) full details of a specific ShapeShift transaction **given the address** *addr_s* in the curIn blockchain (i.e., the address to which the user sent their coins).

数据源

Scrape from exchanger’s public API:

(1) For the 50 most recent transactions, information is provided in the form: (curIn, curOut, amt,t, id)

(2) when provided with a specific **addr_s** ShapeShift provides the tuple (status, **address, withdraw**, inCoin, inType, outCoin, outType, tx, txURL, error)

数据源

Parsed Blockchain Data:

(1) BlockSci for BTC, DASH, ZEC

(2) Python script for remaining 5 cryptocurrencies

Currency	Abbr.	Total	curIn	curOut
Ethereum	ETH	1,385,509	892,971	492,538
Bitcoin	BTC	1,286,772	456,703	830,069
Litecoin	LTC	720,047	459,042	261,005
Bitcoin Cash	BCH	284,514	75,774	208,740
Dogecoin	DOGE	245,255	119,532	125,723
Dash	DASH	187,869	113,272	74,597
Ethereum Classic	ETC	179,998	103,177	76,821
Zcash	ZEC	154,142	111,041	43,101

4

识别跨链交易

(1) **passthrough transactions**, which represent the full flow of money as it moves from one currency to the other via the deposit and withdrawal transactions

(2) **U-turns**, in which a user who has shifted into one currency immediately shifts back

(3) **round-trip transactions**, which are essentially a combination of the first two and follow a user’ s flow of money as it moves from one currency to another and then back to the original one.

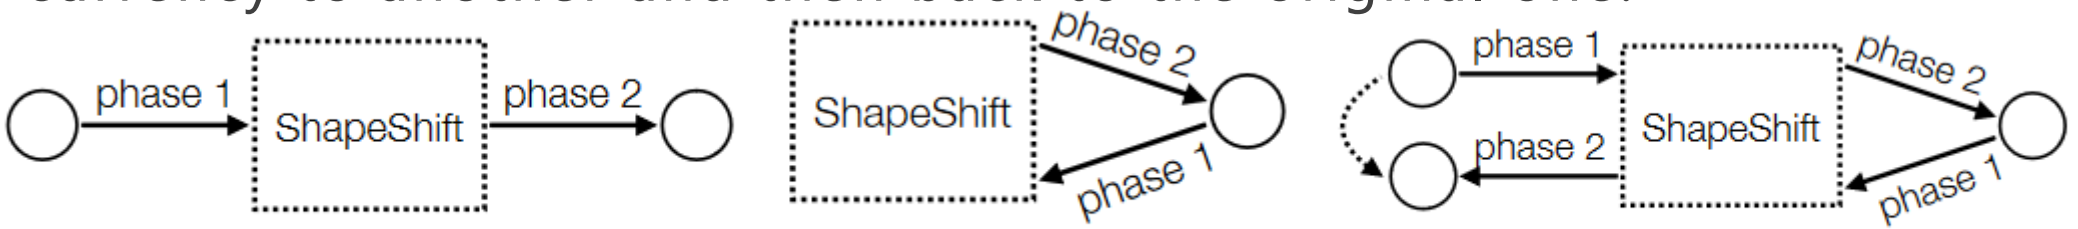

(a) Pass-through (b) U-turn (c) Round-trip

passthrough transactions

Given a ShapeShift transaction with timestamp t, we first find the block b (at some height h) on the *curIn* blockchain that was mined at the time closest to t. We then look at the transactions in all blocks with height in the range $[h - \delta_b, h + \delta_a]$, where δ_b and δ_a are parameters specific to *curIn*.

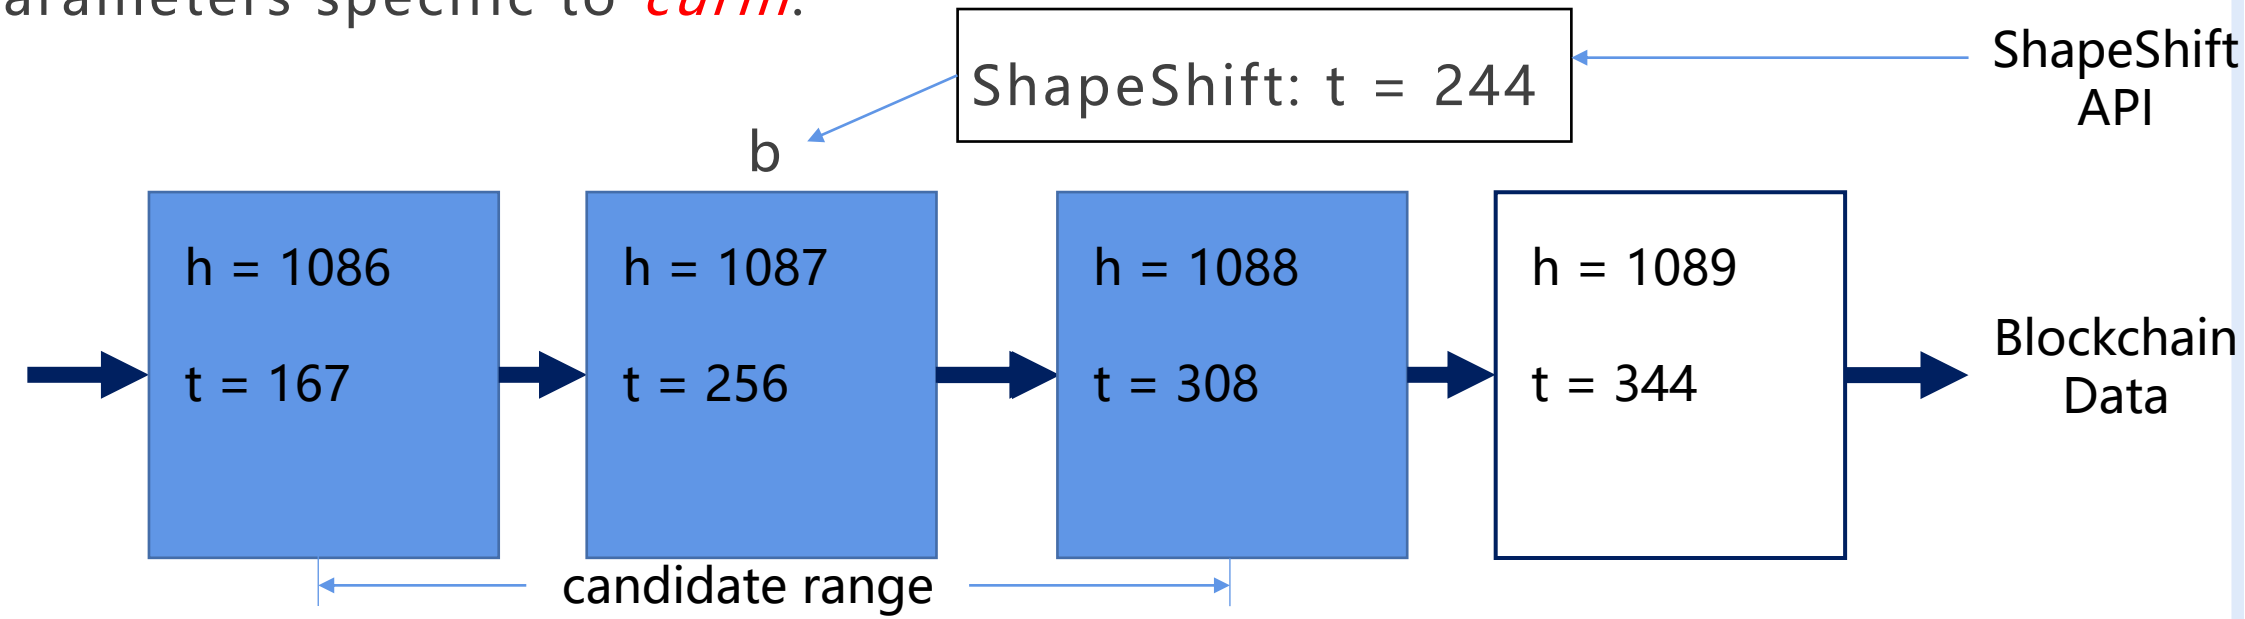

passthrough transactions

we consider two main requirements for identifying the correct on-chain transaction: (1) that it occurred reasonably close in **time** to the point at which it was advertised via the API, and (2) that the **value** it carried was identical to the advertised amount.

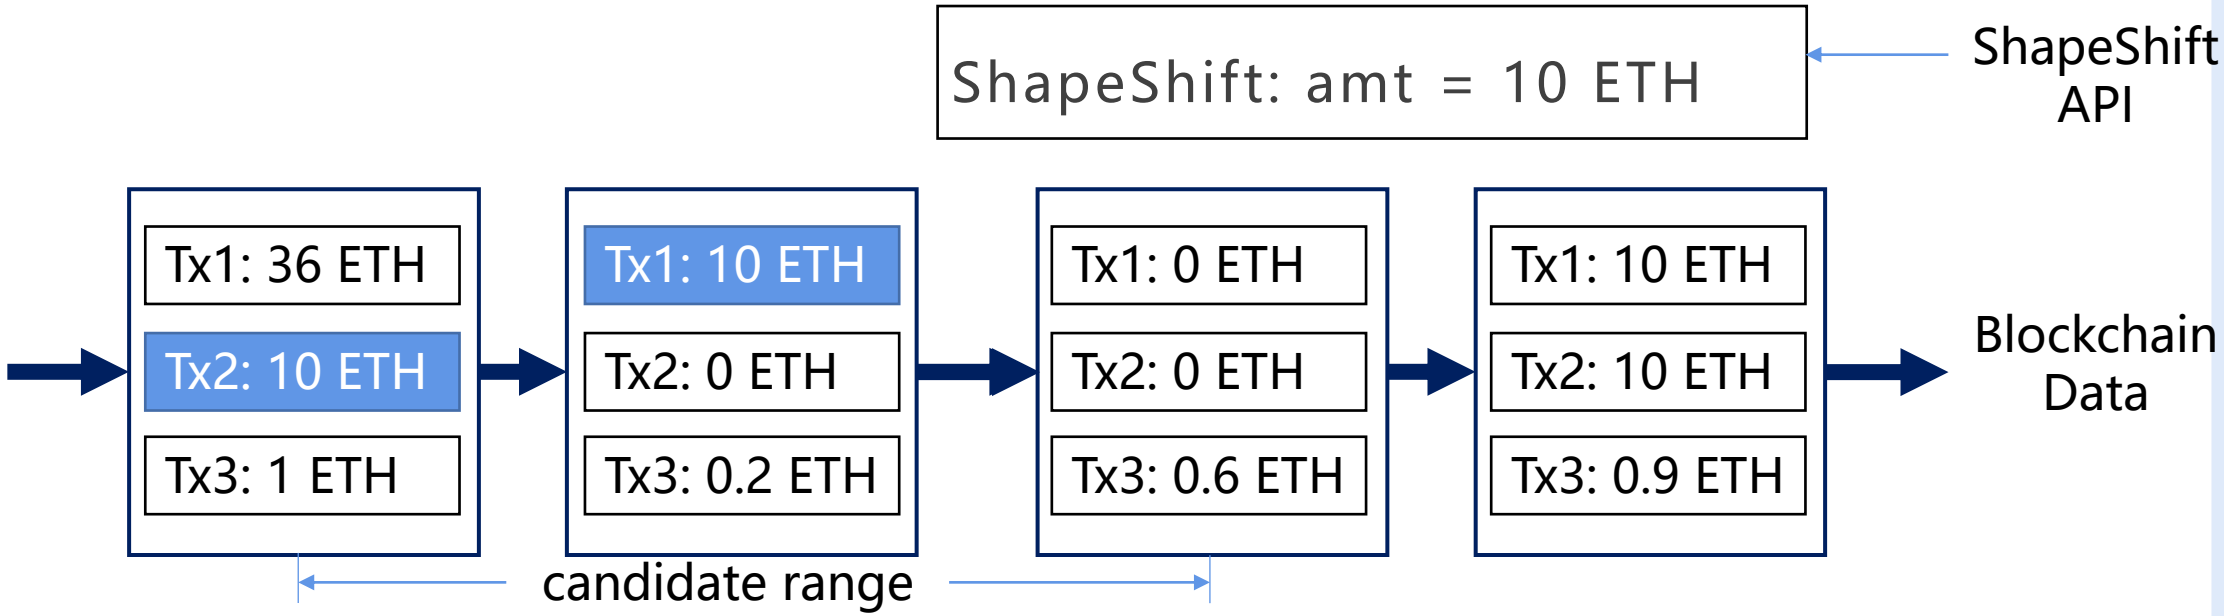

passthrough transactions

we queried the API on the **recipient address** of every transaction identified by our heuristic for Phase 1.

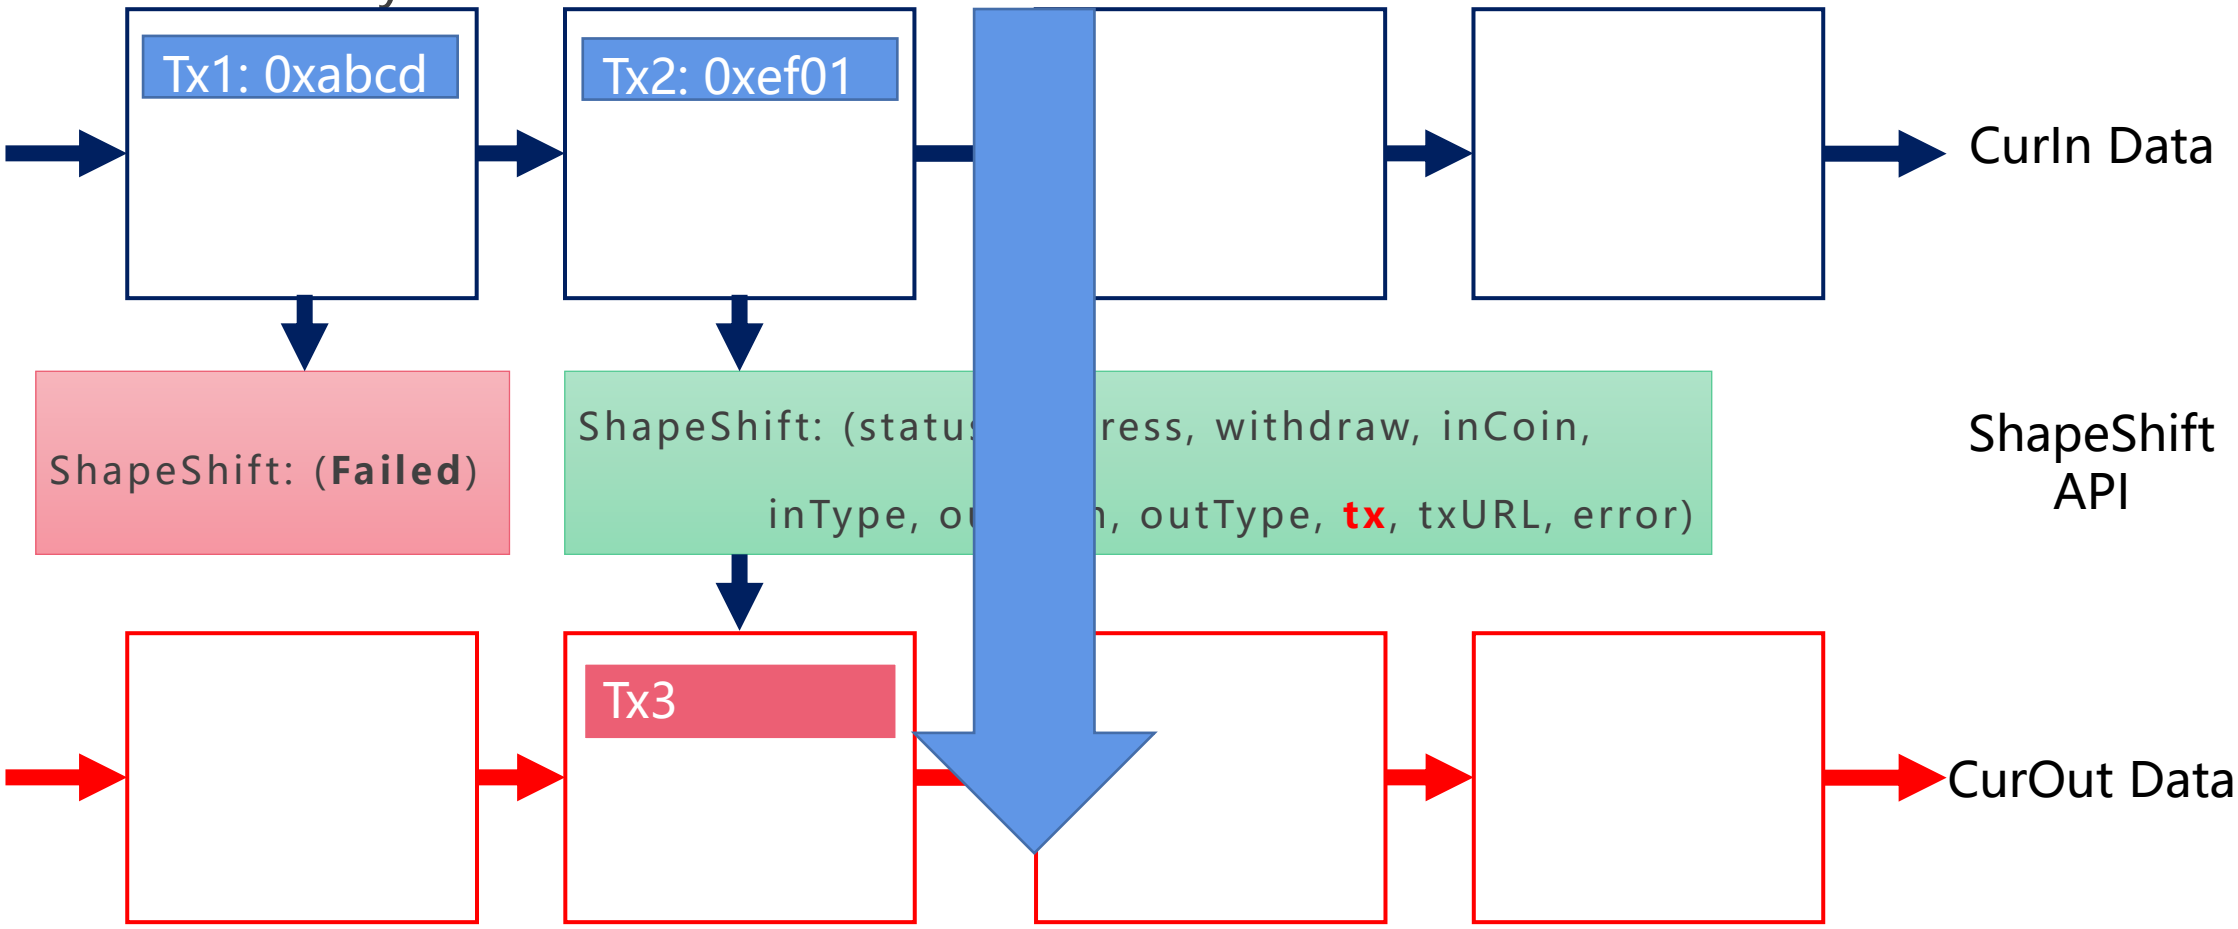

passthrough transactions

We thus ensured that the transaction returned by the API had three things in common with the ShapeShift transaction: (1) the pair of currencies, (2) the amount being sent, and (3) the timing, within the interval specified in Table 2. If there was any mismatch, we discarded the transaction.

Currency	Parameters		Basic %	Augmented %
	δ_b	δ_a		
BTC	0	1	65.76	76.86
BCH	9	4	76.96	80.23
DASH	5	5	84.77	88.65
DOGE	1	4	76.94	81.69
ETH	5	0	72.15	81.63
ETC	5	0	76.61	78.67
LTC	1	2	71.61	76.97
ZEC	1	3	86.94	90.54

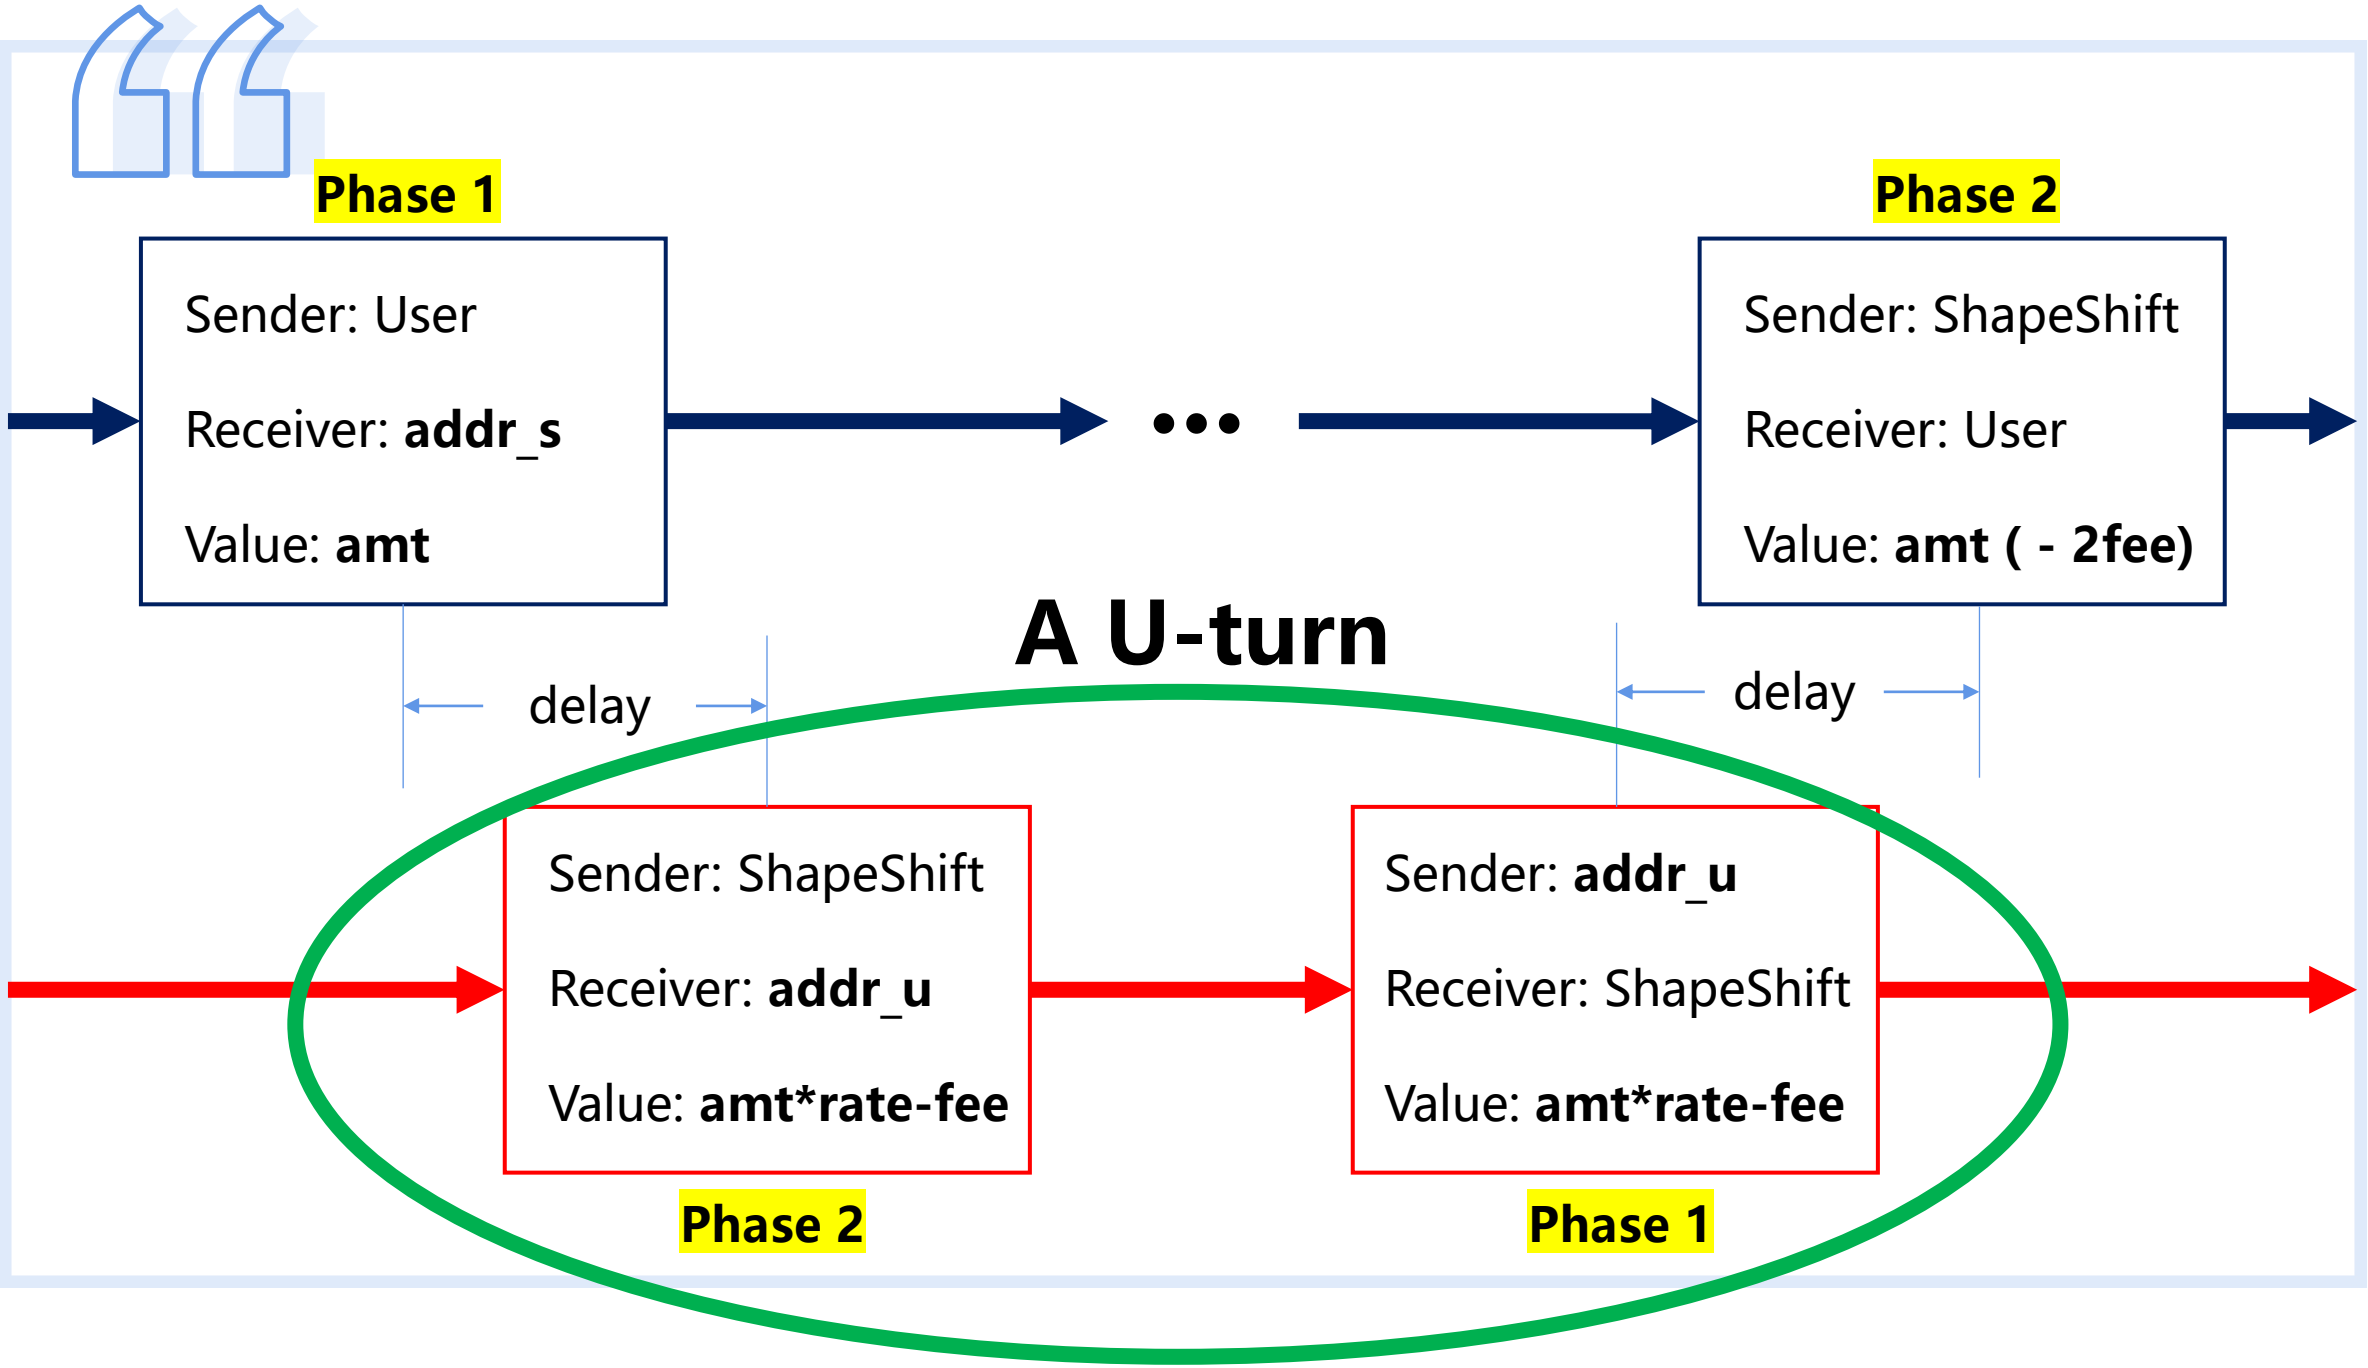
Phase 1
Sender: User
Receiver: addr_s
Value: amt
Phase 2
Sender: ShapeShift
Receiver: User
Value: amt (- 2fee)
A U-turn
delay
delay
Sender: ShapeShift
Receiver: addr_u
Value: amt*rate-fee
Sender: addr_u
Receiver: ShapeShift
Value: amt*rate-fee
Phase 2
Phase 1

U-turn transactions

1.Basic timing&value based heuristic

we require that the second transaction happens within **30 minutes** of the first, and that it carries within **1%** of the value that was generated by the first Phase 2 transaction.

2.UTXO-based heuristic

we could see if the input **is the same UTXO** that was created in the Phase 2 transaction

3.Account-based heuristic

Here we thus define a U-turn as seeing if **the address that was used as the output** in the Phase 2 transaction **is used as the input** in the later Phase 1 transaction.

U-turn transactions

Currency	# (basic)	# (addr)	# (utxo)
BTC	36,666	565	314
BCH	2864	196	81
DASH	3234	2091	184
DOGE	546	75	75
ETH	53,518	5248	-
ETC	1397	543	-
LTC	8270	1429	244
ZEC	772	419	222

round-trip transactions

(1)the first transaction was of the form X-Y;

(2)the second transaction was of the form Y-X;

(3)the second transaction happened relatively soon after the first one;

(4)the value carried by the two transaction was approximately the same.

For the third property, we required that the second transaction happened within **30 minutes** of the first. For the fourth property, we required that if the first transaction carried x units of curIn then the second transaction carried within **0.5% of the value** in the (on-chain) Phase 2 transaction, according to the outCoin field provided by the API.

round-trip transactions

Currency	# (regular)	# (same addr)
BTC	35,019	437
BCH	1780	84
DASH	3253	2353
DOGE	378	0
ETH	45,611	4085
ETC	1122	626
LTC	6912	2733
ZEC	472	172

Table 4: The number of regular round-trip transactions identified for each cryptocurrency, and the number that use the same initial and final address.

5

图分析

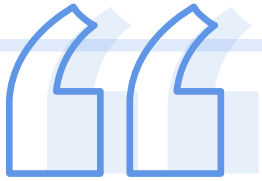

Heuristic 7.1. *If two or more addresses send coins to the same address in the* curOut *blockchain, or if two or more addresses receive coins from the same address in the* curIn *blockchain, then these addresses have* ***some common*** *social relationship.*

To implement this heuristic, we parsed transactions into a graph where **we defined a node as an address** and **a directed edge (u, v)** as existing when **one address u initiated a ShapeShift transaction sending coins to v**. Edges are further **weighted by the number of transactions** sent from u to v. For each node, the cluster centered on that address was then defined as all nodes adjacent to it.

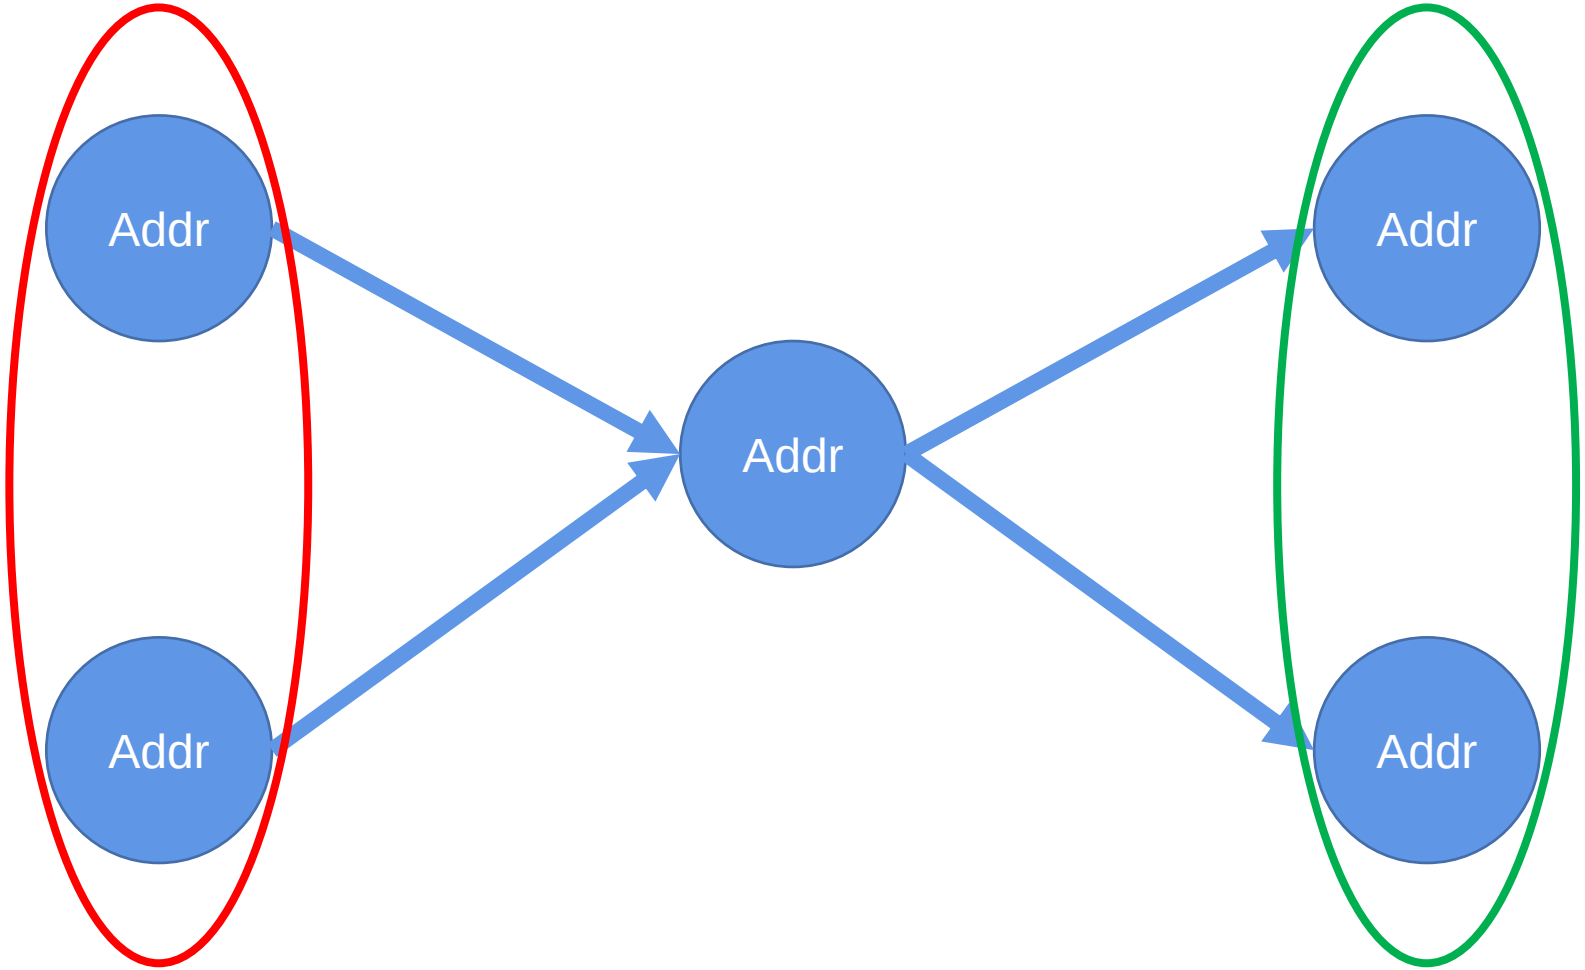
Addr
Addr
Addr
Addr
Addr
the input cluster
the output cluster

Sorting the clusters **by in-degree** reveals the entities that received the highest number of ShapeShift transactions (from the top 8 currencies, per our caveat above). The largest cluster had 12,868 addresses — many of them belonging to Ethereum, Litecoin, and Dash — and was centered on a Bitcoin address belonging to CoinPayments.net, a multi-coin payment processing gateway.

Sorting the individual clusters **by out-degree** reveals instead the users who initiated the highest number of ShapeShift transactions. Here the largest cluster (consisting of 2314 addresses) was centered on a Litecoin address, and the second largest cluster was centered on an Ethereum address that belonged to Binance (a popular exchange).

6

应用实例

1.Zcash

We considered three possible interactions between ShapeShift and the shielded pool, as depicted in Figure: (1) a user shifts coins directly from ShapeShift into the shielded pool, (2) a user shifts to a t-address but then uses that t-address to put money into the pool, and (3) a user sends money directly from the pool to ShapeShift.

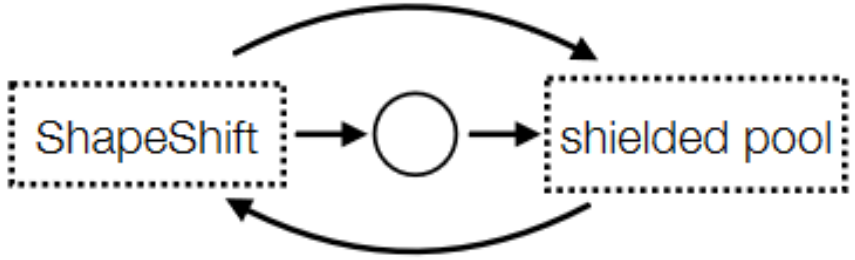

Figure 7: The three types of interactions we investigated between ShapeShift and the shielded pool in Zcash.

1.ZCash

For **the first type** of interaction, we found 29,003 transactions that used ZEC as curOut. Of these, 758 had a z-address as the output address, meaning coins were **sent directly to** the shielded pool. The total value put into the pool in these transactions was 6,707.86 ZEC, which is 4.3% of all the ZEC received in pass-through transactions.

For **the second type** of interaction, there were 1309 where **the next** transaction (i.e., the transaction in which this UTXO spent its contents) involved putting money into the pool. The total value put into the pool in these transactions was 12,534 ZEC, which is 8.2% of all the ZEC received in passthrough transactions.

1.Zcash

For **the third type** of interaction, we found 111,041 passthrough transactions that used ZEC as curIn. Of these, 3808 came directly from the pool, with a total value of 22,490 ZEC (14% of all the ZEC sent in pass-through transactions).

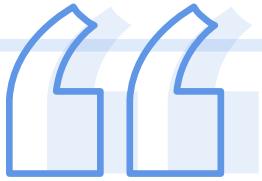

2.Dash

Our parameters for identifying a CoinJoin were thus that (1) the transaction must have **at least three inputs**, (2) the outputs must consist solely of values from the list of possible denominations (**modulo the fees**), and (3) and all **output values must be the same**. In fact, given how Dash operates there is always one output with a non-standard value, so it was further necessary to relax the second and third requirements to allow there to be at most one address that does not carry the specified value.

2.Dash

We first looked to identify if **the inputs of a Phase 1** transaction were outputs from a CoinJoin. Out of 100,410 candidate transactions, we found 2,068 that came from a CoinJoin, carrying a total of 11,929 DASH in value (**6.5%** of the total value across transactions with Dash as curIn). Next, we looked to identify if **the outputs of a Phase 2** transaction had been spent in a CoinJoin. Out of 50,545 candidate transactions, we found only 33 CoinJoin transactions, carrying a total of 187 DASH in value (**0.1%** of the total value across transactions using Dash as curOut).

2.Dash

If we revisit our results concerning the use of U-turns in Dash from Section 6.2, we recall that there was a large asymmetry in terms of the results of our two heuristics: only 5.6% of the **U-turns** used the same UTXO, but 64.6% of U-turns used the same address. This suggests that some additional on-chain transaction took place between the two ShapeShift transactions, and indeed upon further inspection we identified many cases where this transaction was a CoinJoin. There thus appears to have been a genuine attempt to take advantage of the privacy that Dash offers, but this was **completely ineffective** due to the use of the same address that both sent and received the mixed coins.