Advances in Elliptic-Curve Cryptography

Tanja Lange

Technische Universiteit Eindhoven

tanja@hyperelliptic.org

24.08.2011

ECDSA I

- Elliptic-Curve Digital Signature Algorithm.
- Suggested by Certicom, standardized by FIPS and others.
- Some other elliptic-curve signature schemes were proposed, e.g. by Korea and by Germany.
- Starting point is an elliptic curve E, in Weierstrass form over a finite field $\mathbb{F}_q$ and a point $P \in E(\mathbb{F}_q)$ of some prime order ℓ.
- Assume for simplicity $q > 3$ is an odd prime,

$$E : y^2 = x^3 + a_4 x + a_6;\ a_4, a_6 \in \mathbb{F}_q$$

and $E(\mathbb{F}_q) = \langle P \rangle$.

ECDSA II

- Alice's secret key is integer $a \in [1, \ell - 1]$ and her public key is $P_A = aP$.
- To sign a message m with hash value $h =$hash(m) she
  - picks random $k \in [1, \ell - 1]$;
  - computes $R = kP$ and puts $r \equiv x(R) \bmod \ell$;
  - computes $s \equiv k^{-1}(h + ra) \bmod \ell$.
  - The signature is the tuple (r, s), both values are smaller than ℓ.
- To verify the signature (r, s) on m Bob computes
  - $w \equiv s^{-1} \bmod \ell, u_1 \equiv hw \bmod \ell, u_2 \equiv rw \bmod \ell$;
  - and $Q = u_1 P + u_2 P_A$;
  - and accepts the signature if $r \equiv x(Q) \bmod \ell$; otherwise he rejects it.

ECDSA III

- Alice's signatures should be accepted, i.e. the x-coordinate of $R = kP$ should be the same as the x-coordinate of $Q = u_1P + u_2P_A$, where $P_A = aP$, $w \equiv s^{-1} \bmod \ell$, $u_1 \equiv hw \bmod \ell$, $u_2 \equiv rw \bmod \ell$, $r \equiv x(R) \bmod \ell$, and $s \equiv k^{-1}(h + ra) \bmod \ell$.

$$\begin{aligned} Q &= u_1P + u_2P_A = hwP + rwP_A = s^{-1}(hP + raP) \\ &= (k^{-1}(h + ra))^{-1}(h + ra)P = kP = R \end{aligned}$$

- Note that P having order ℓ takes care of $\bmod \ell$ computations. Since $\ell \approx q$, $\pm R$ essentially given by r
- Only secret part in (r, s) is a; attacker Eve can choose any random k.
- Computing a from P and $P_A = aP$ is solving the Discrete Logarithm Problem (DLP) of P_A with base P.

The rho method

- Standard attack on DLP: parallel rho.
- Expect $(1+o(1))\sqrt{\pi\ell/2}$ group operations, matching Nechaev/Shoup bound.
- Easy to distribute across CPUs; very little memory consumption; very little communication.
- Simplified, non-parallel rho: Make a pseudo-random walk in $\langle P\rangle$, where the next step depends on current point: $W_{i+1} = f(W_i)$.
- Birthday paradox: Randomly choosing from ℓ elements picks one element twice after about $\sqrt{\pi\ell/2}$ draws.
- The walk now enters a cycle. Cycle-finding algorithm (e.g., Floyd) quickly detects this.

Random walk example

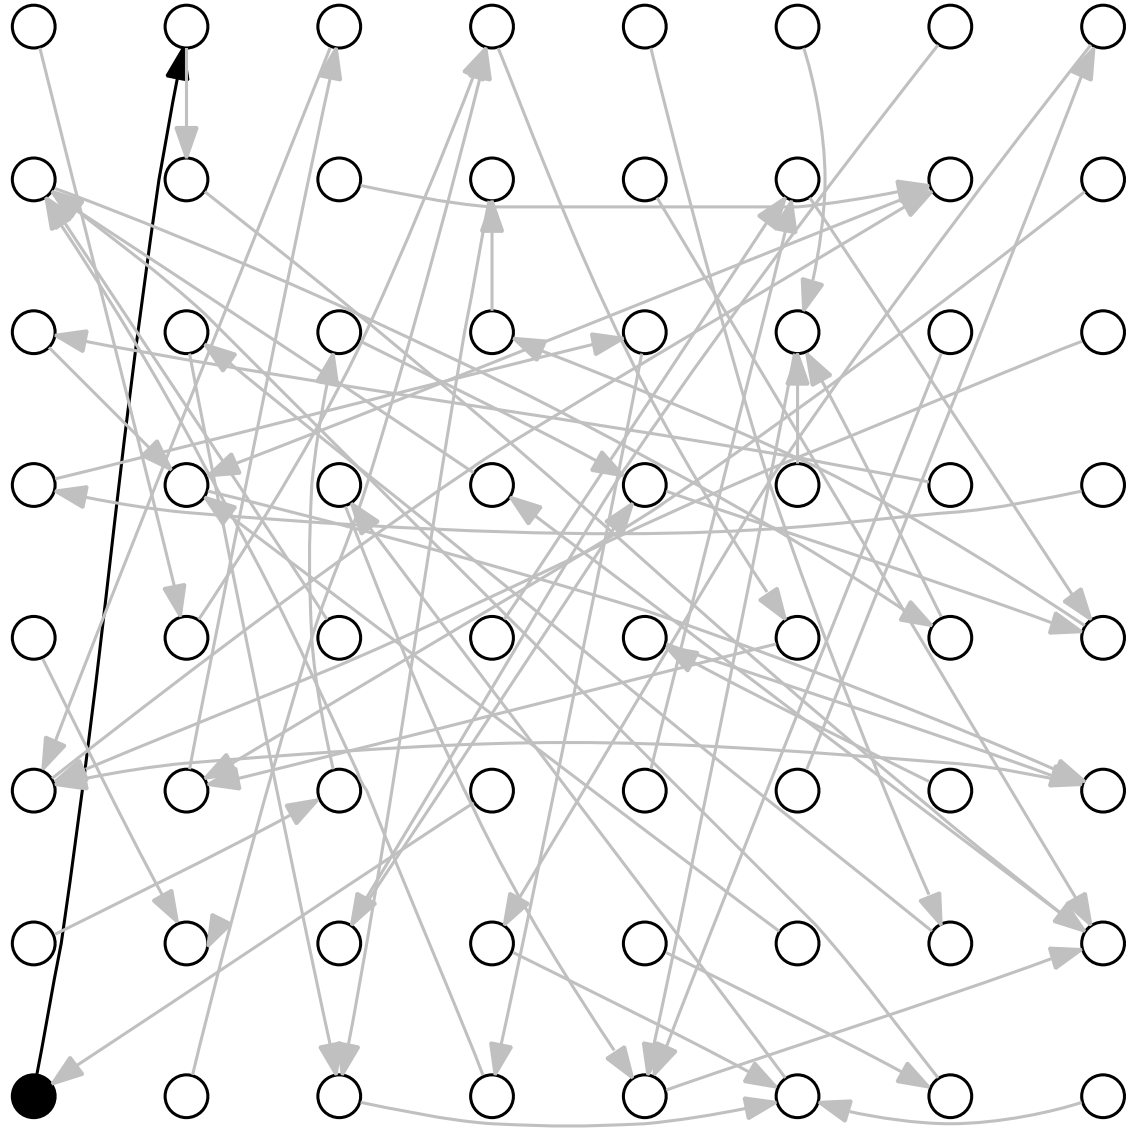

Random walk example

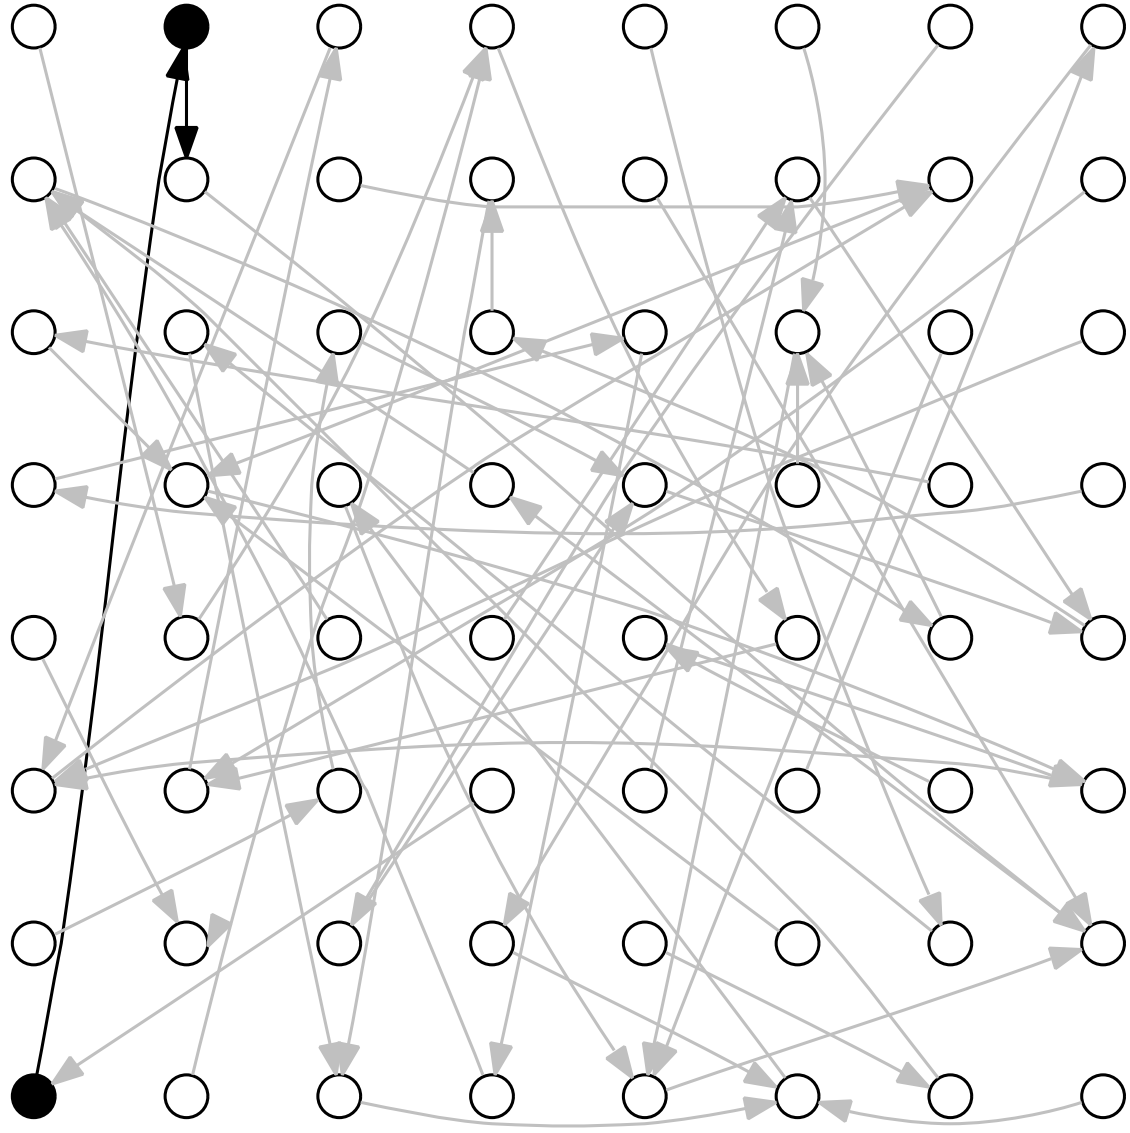

Random walk example

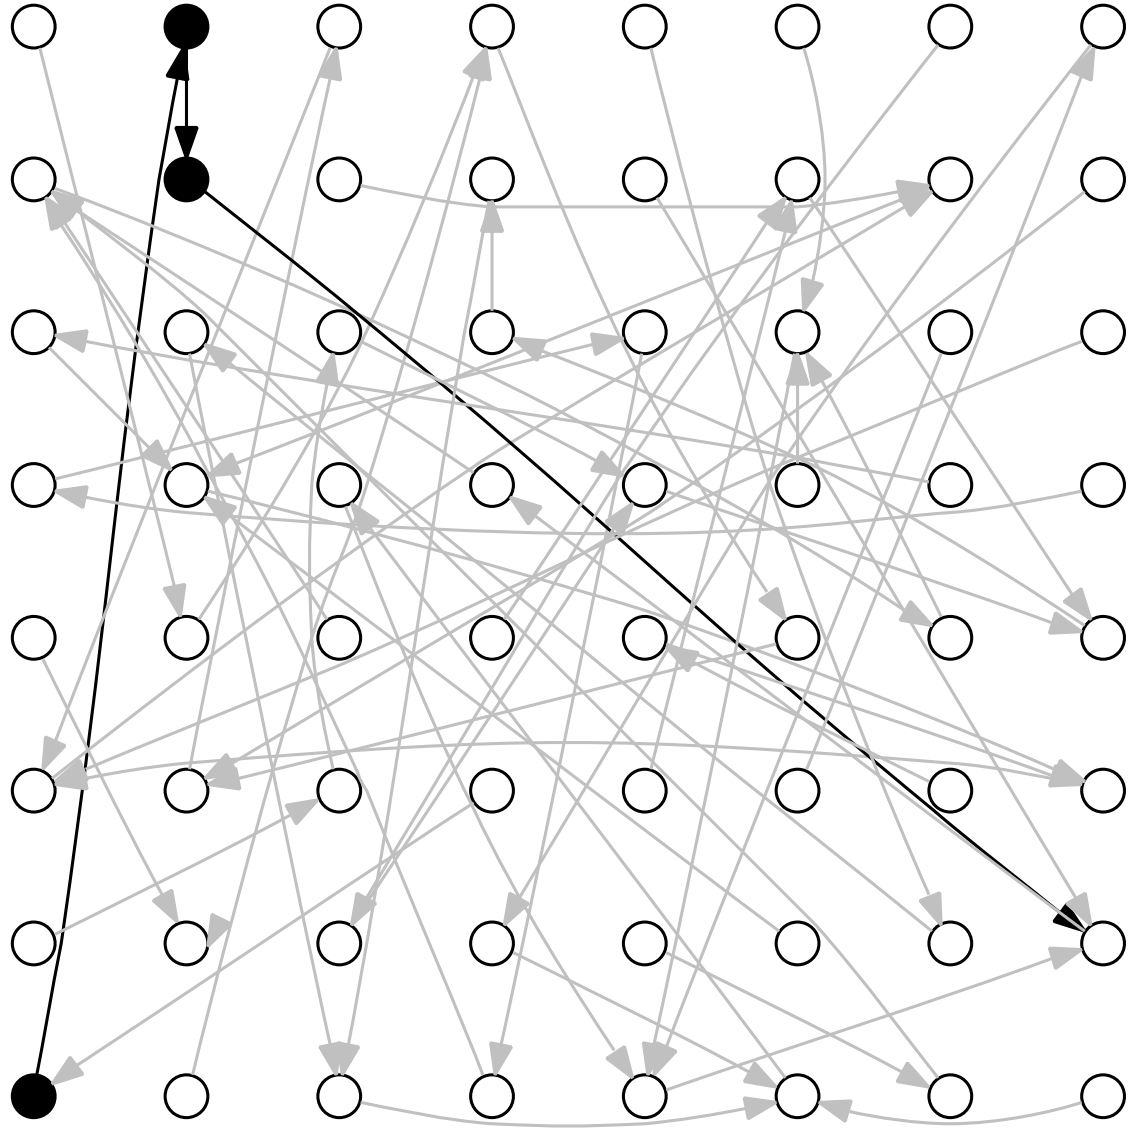

Random walk example

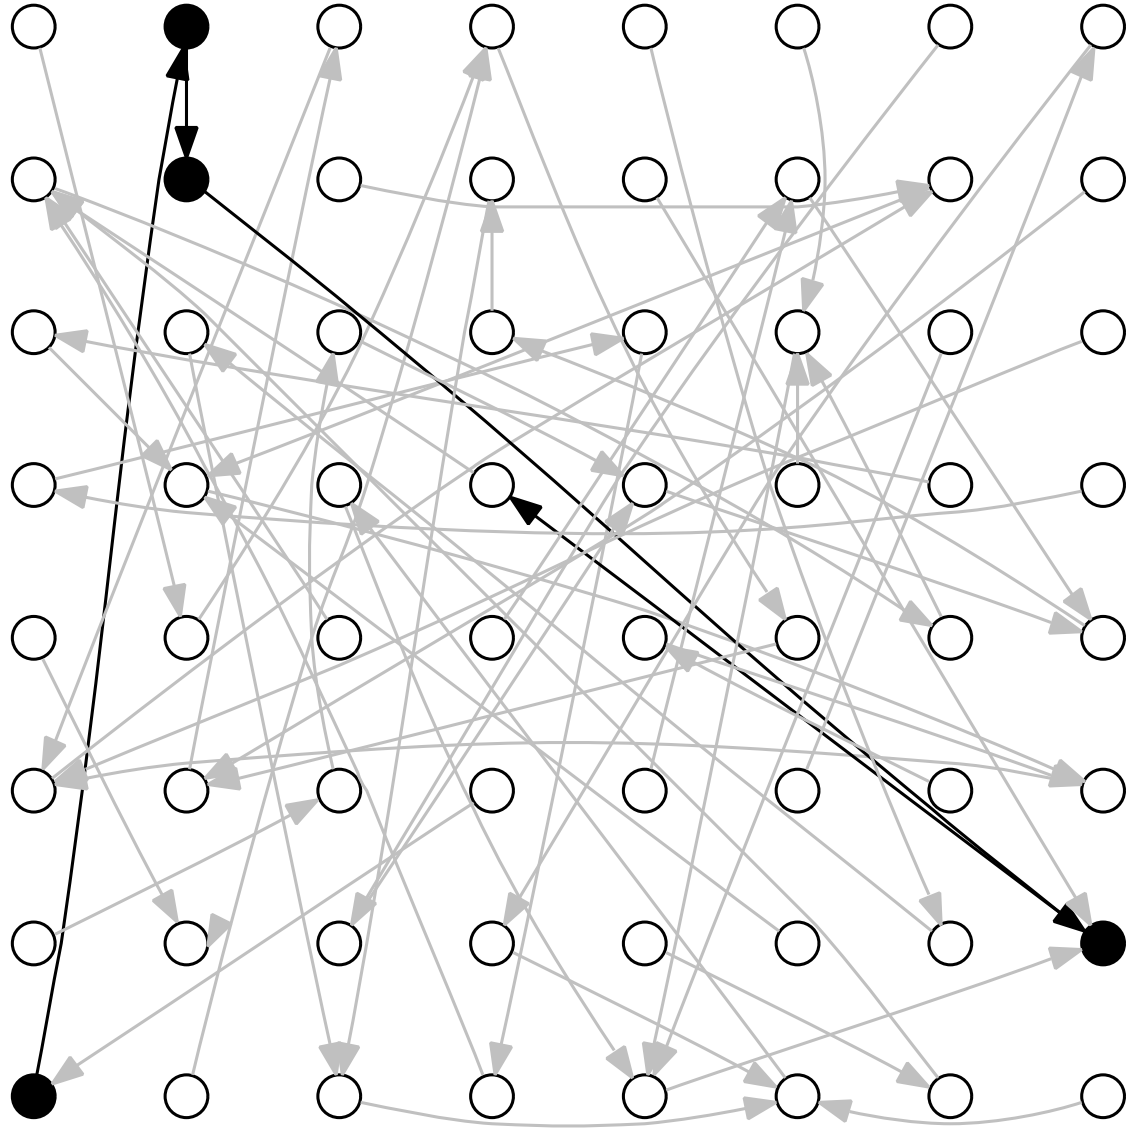

Random walk example

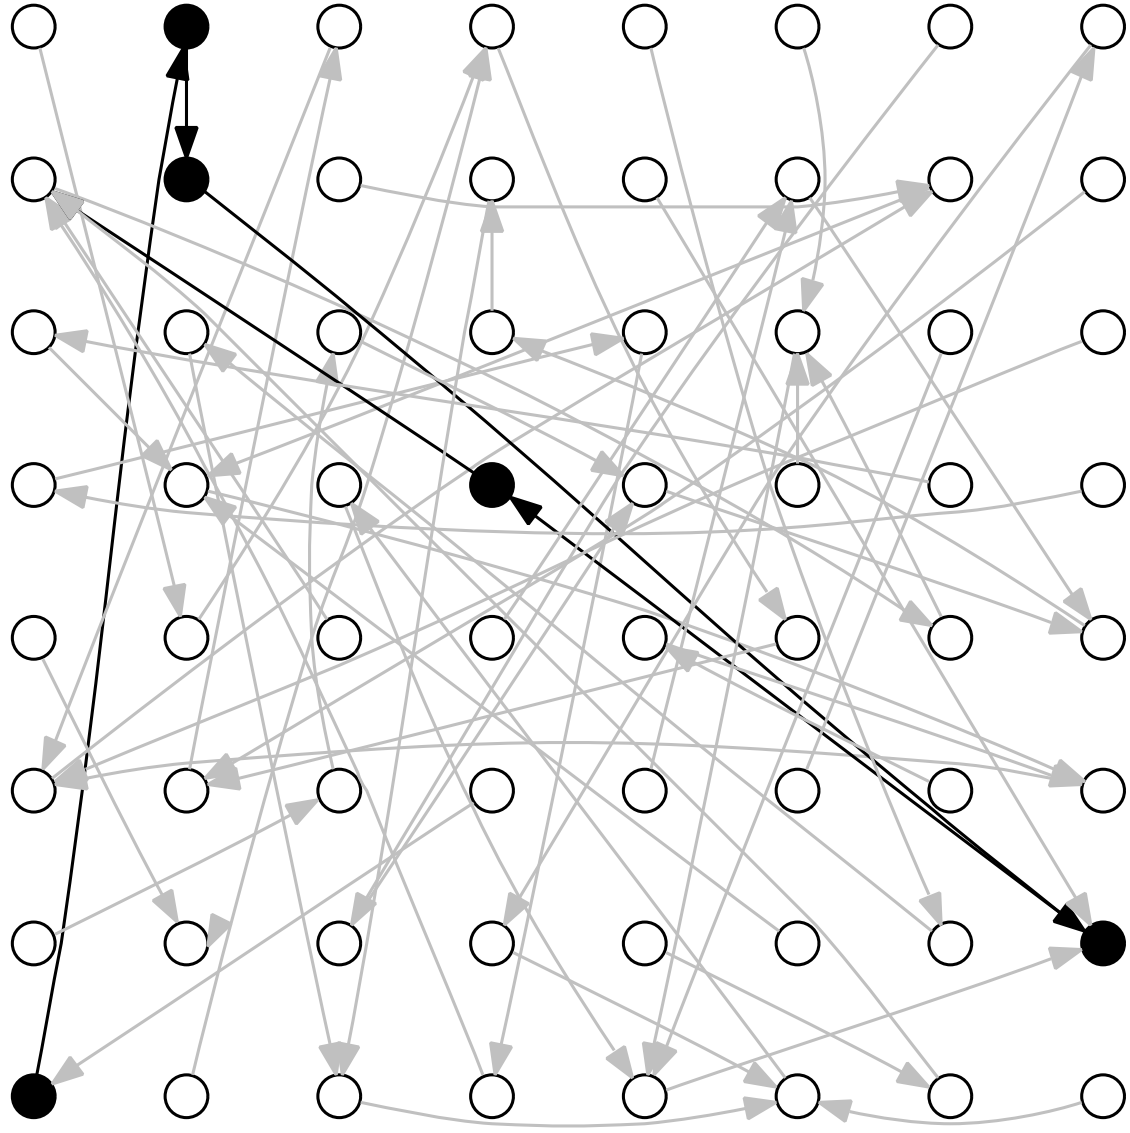

Random walk example

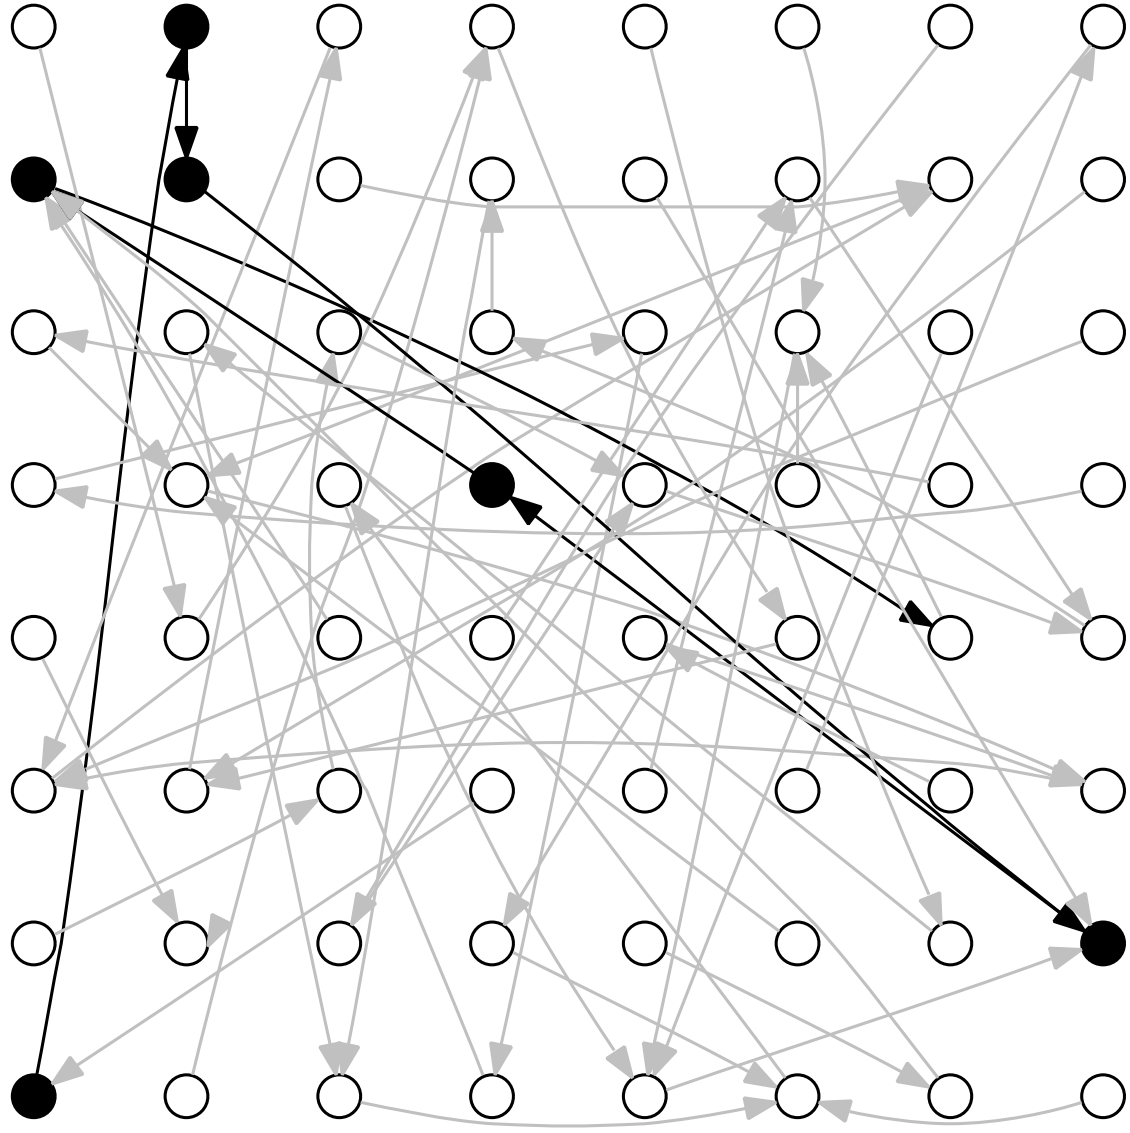

Random walk example

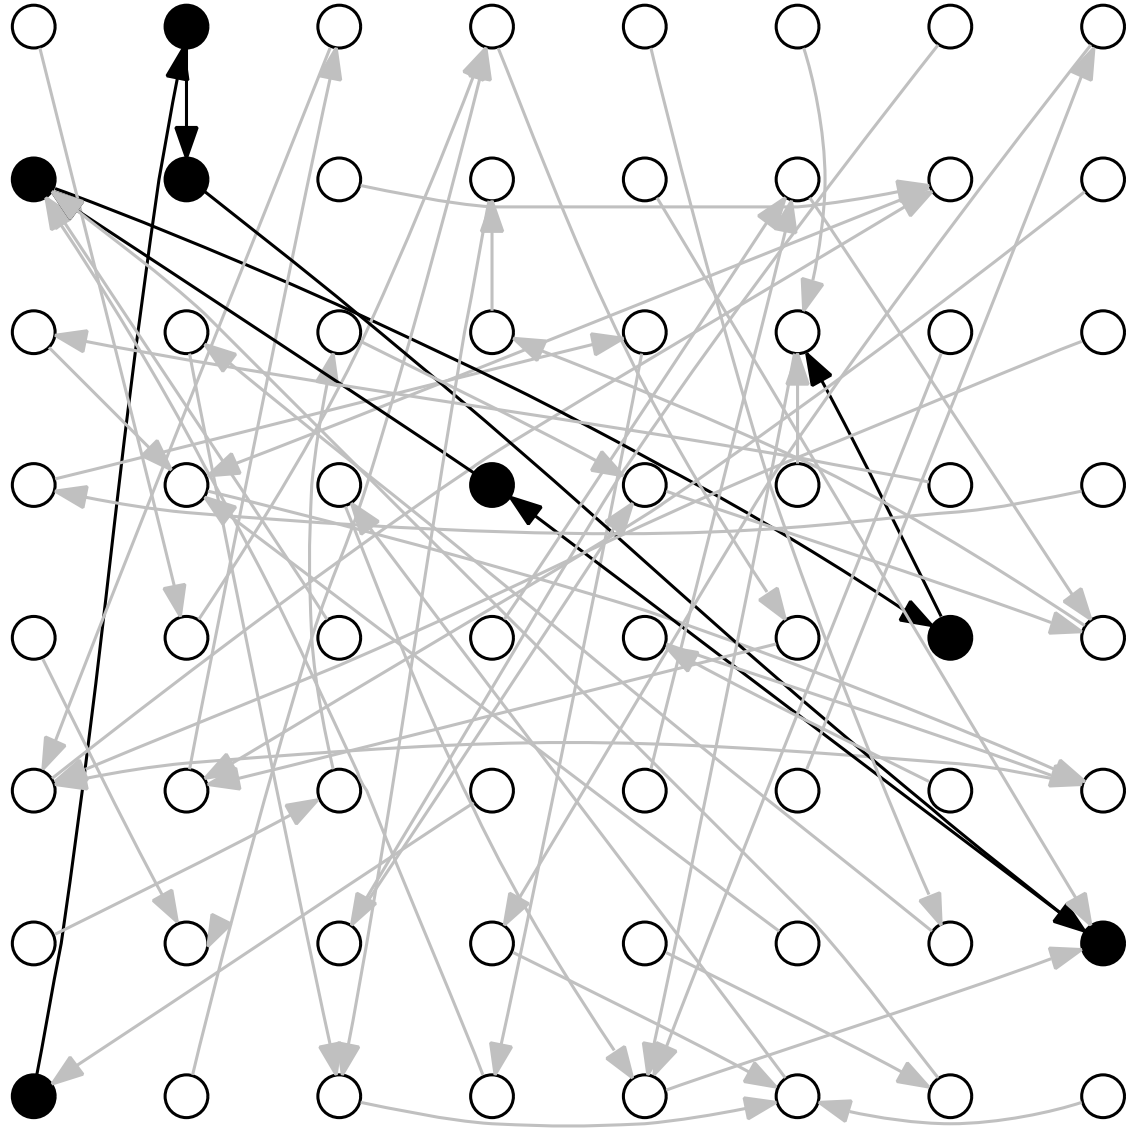

Random walk example

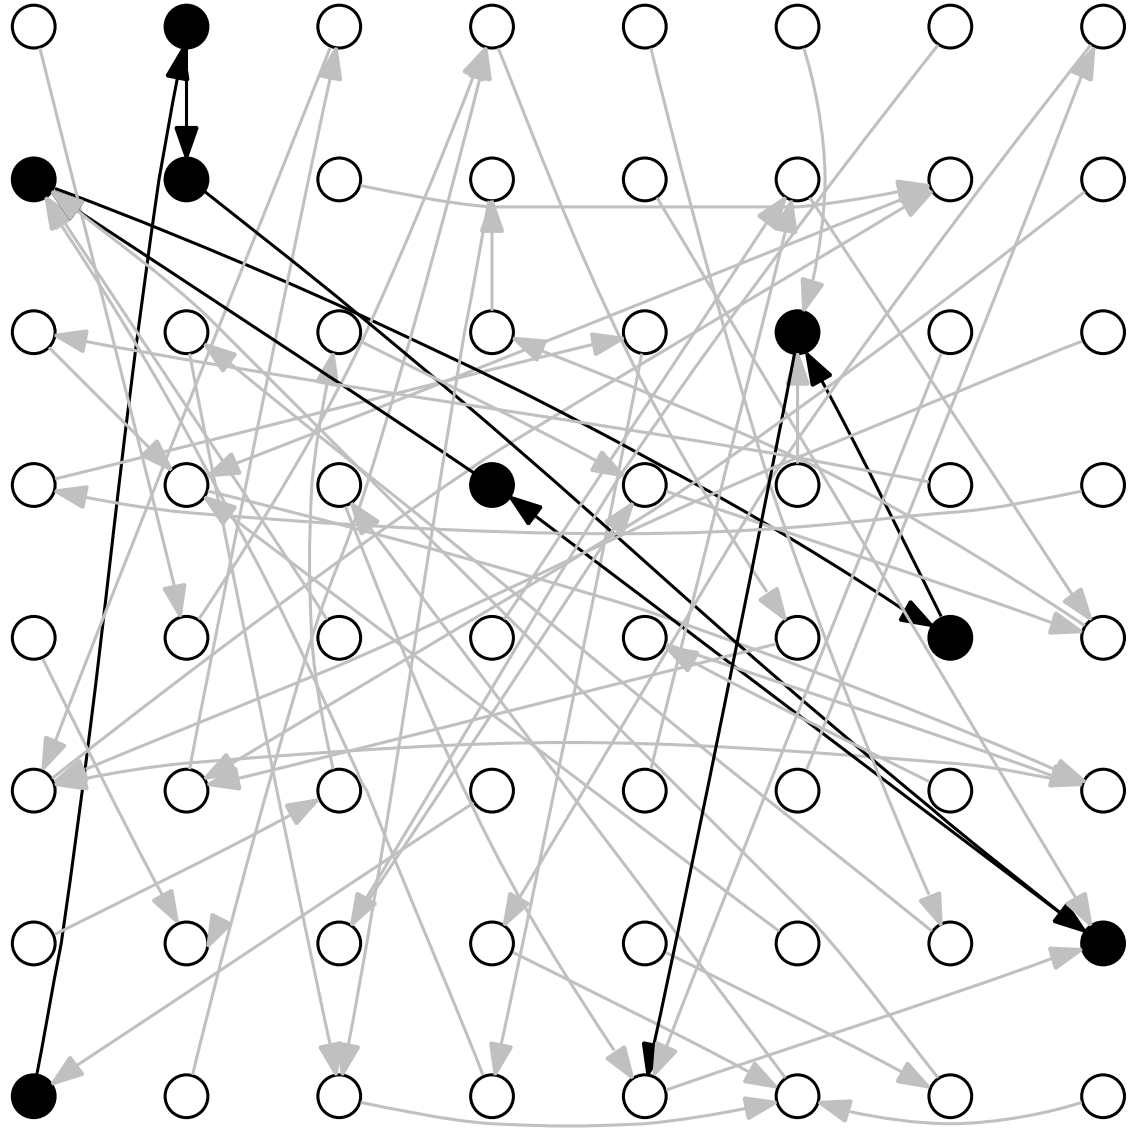

Random walk example

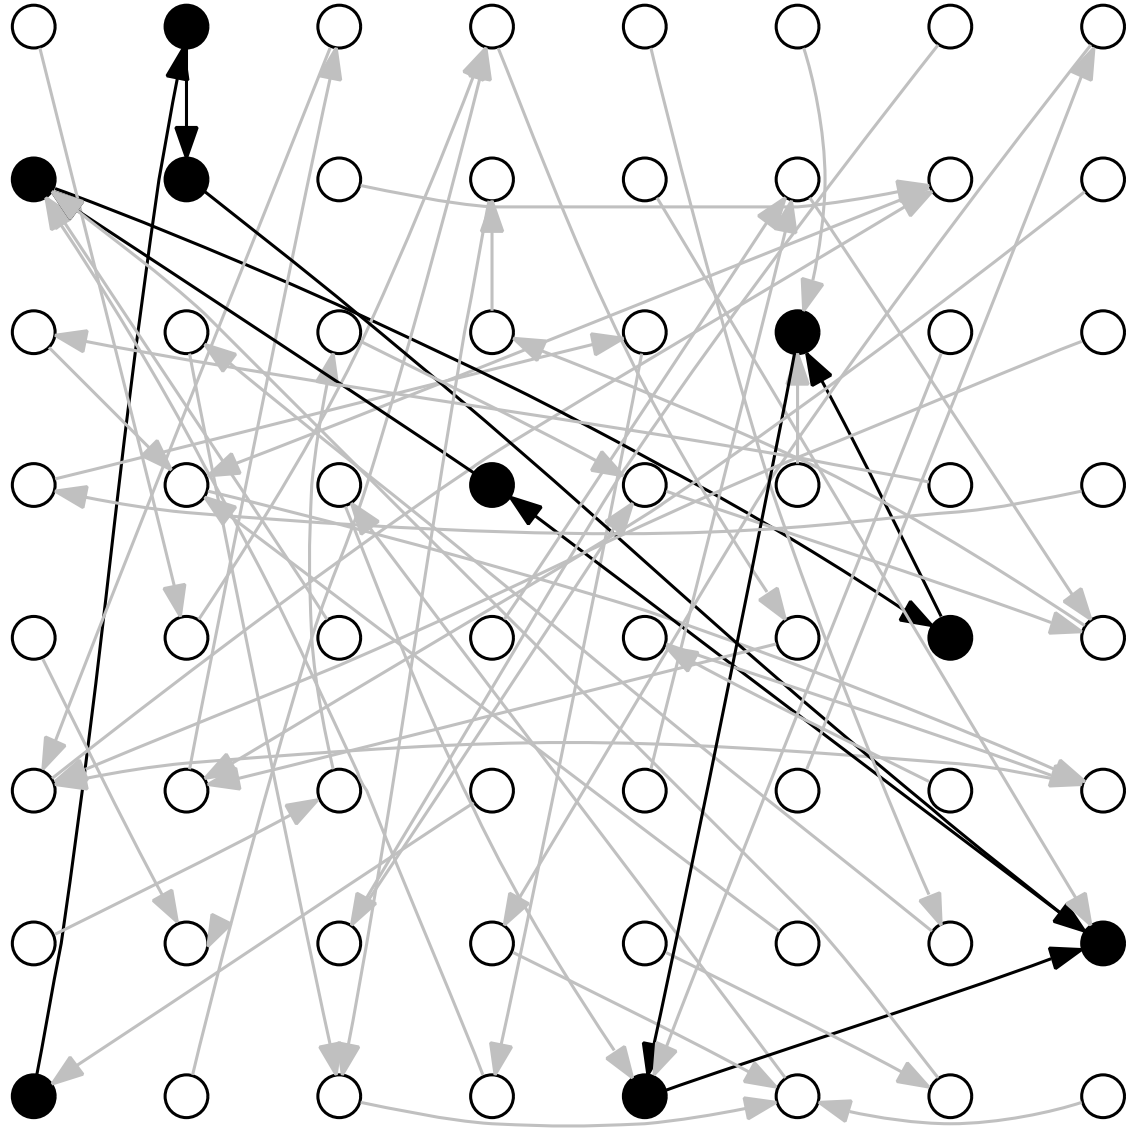

Random walk example

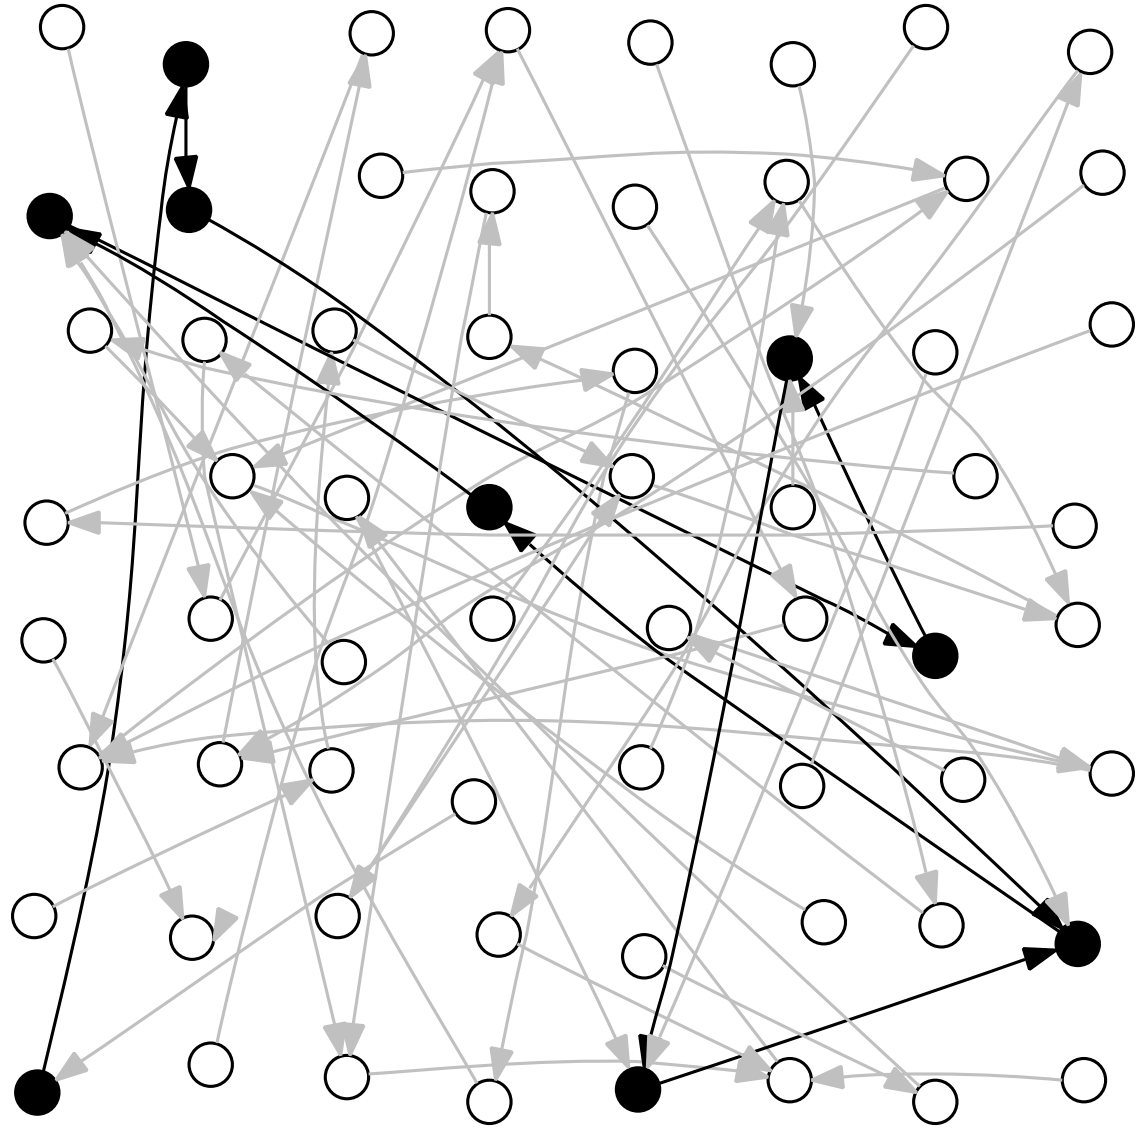

Random walk example

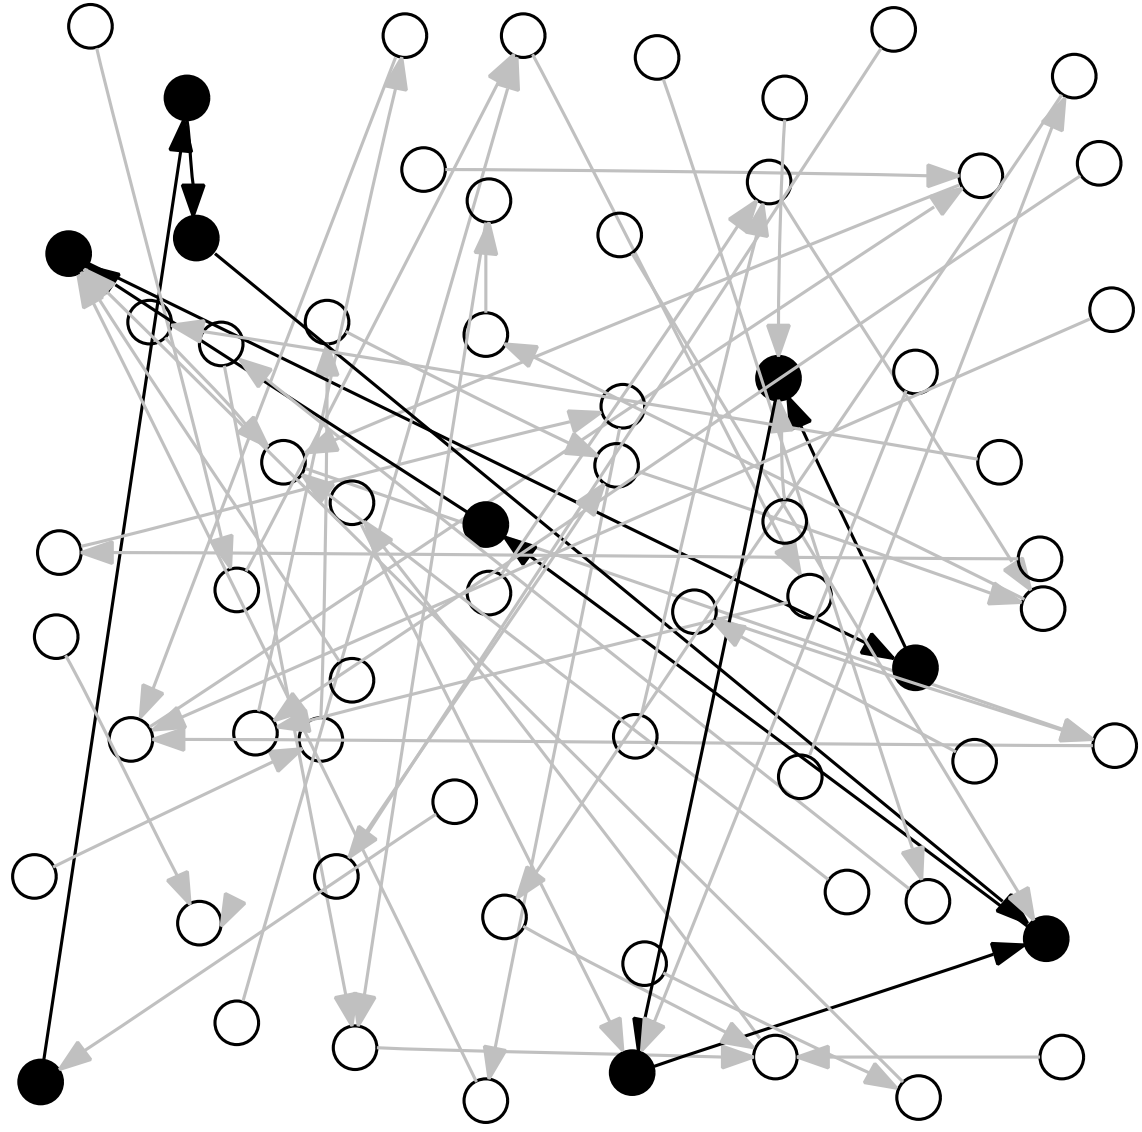

Random walk example

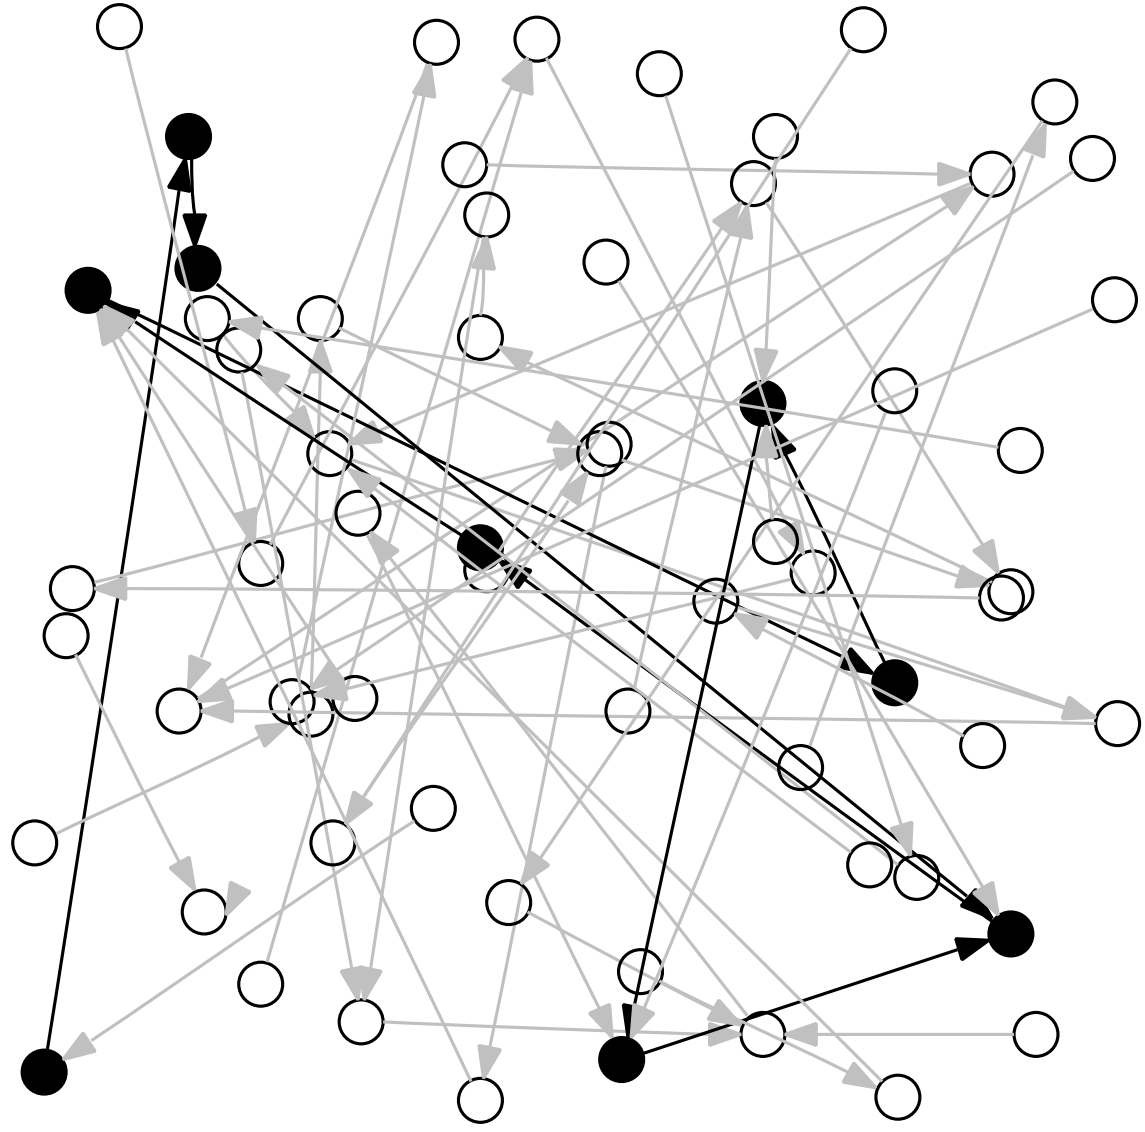

Random walk example

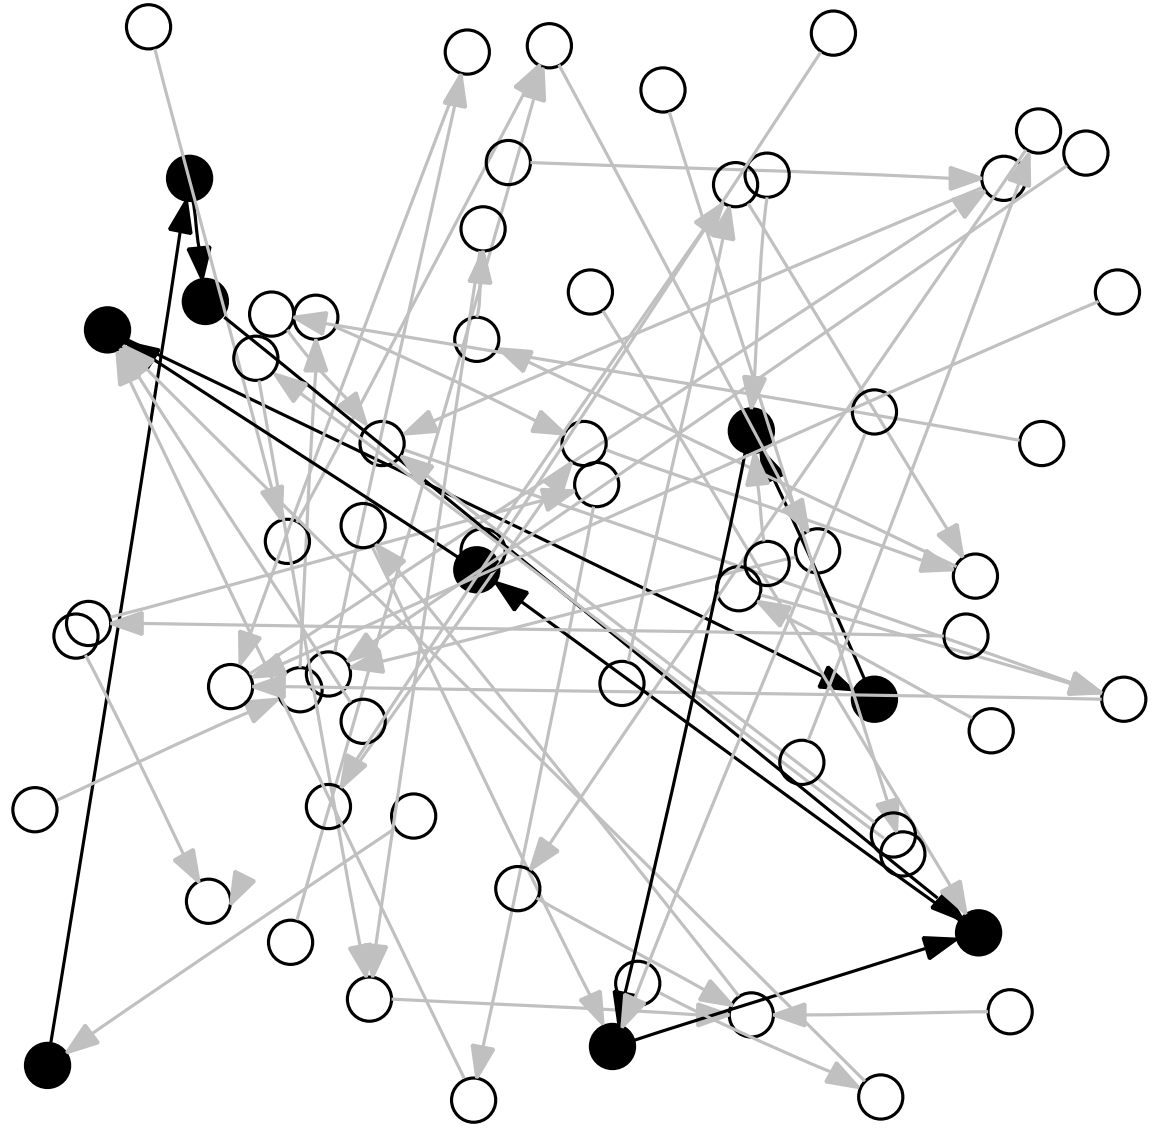

Random walk example

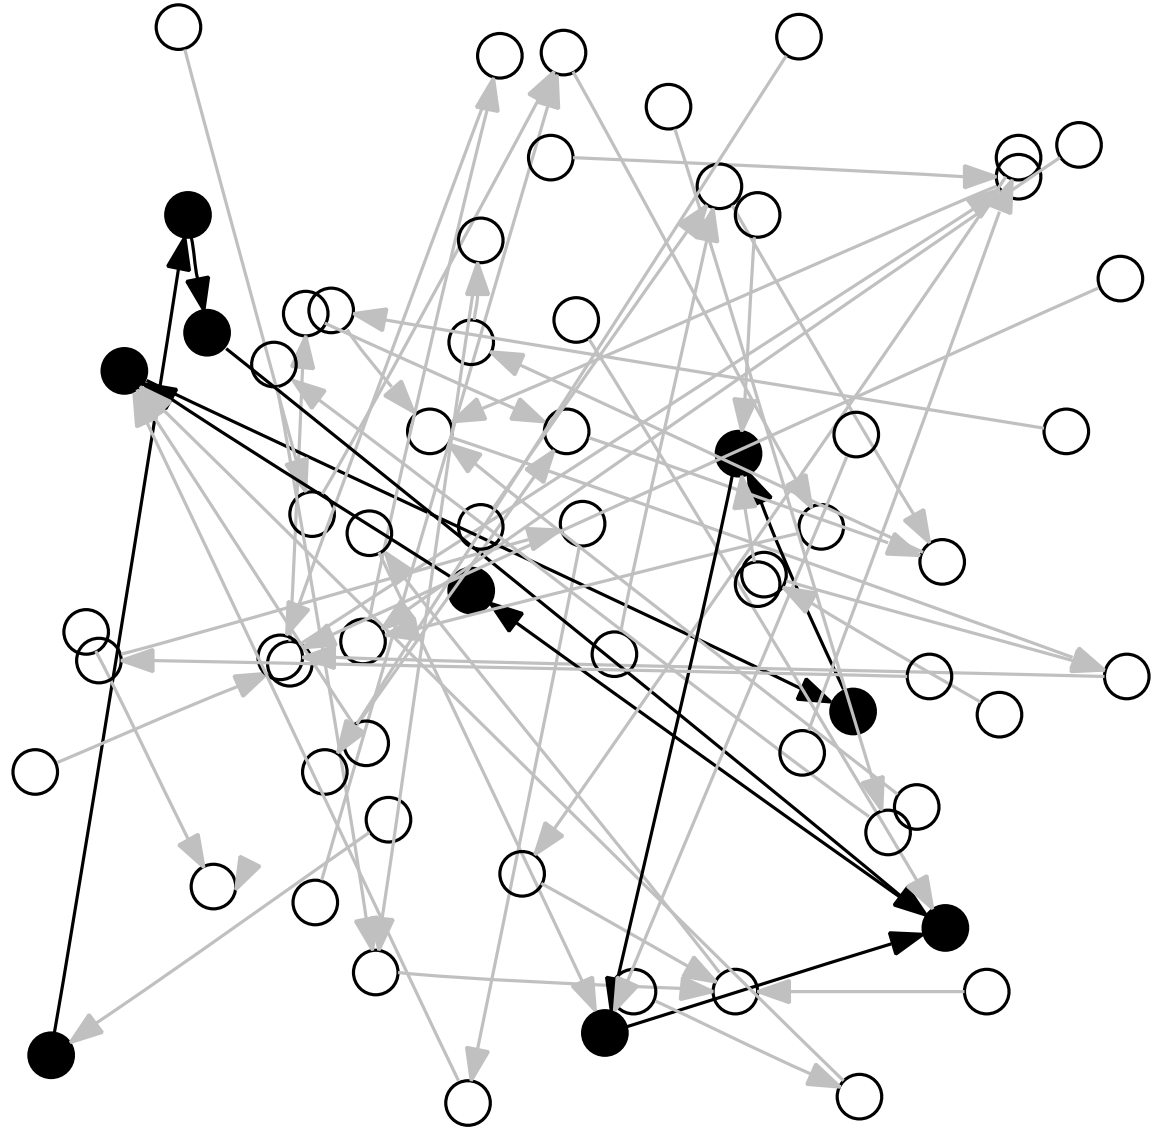

Random walk example

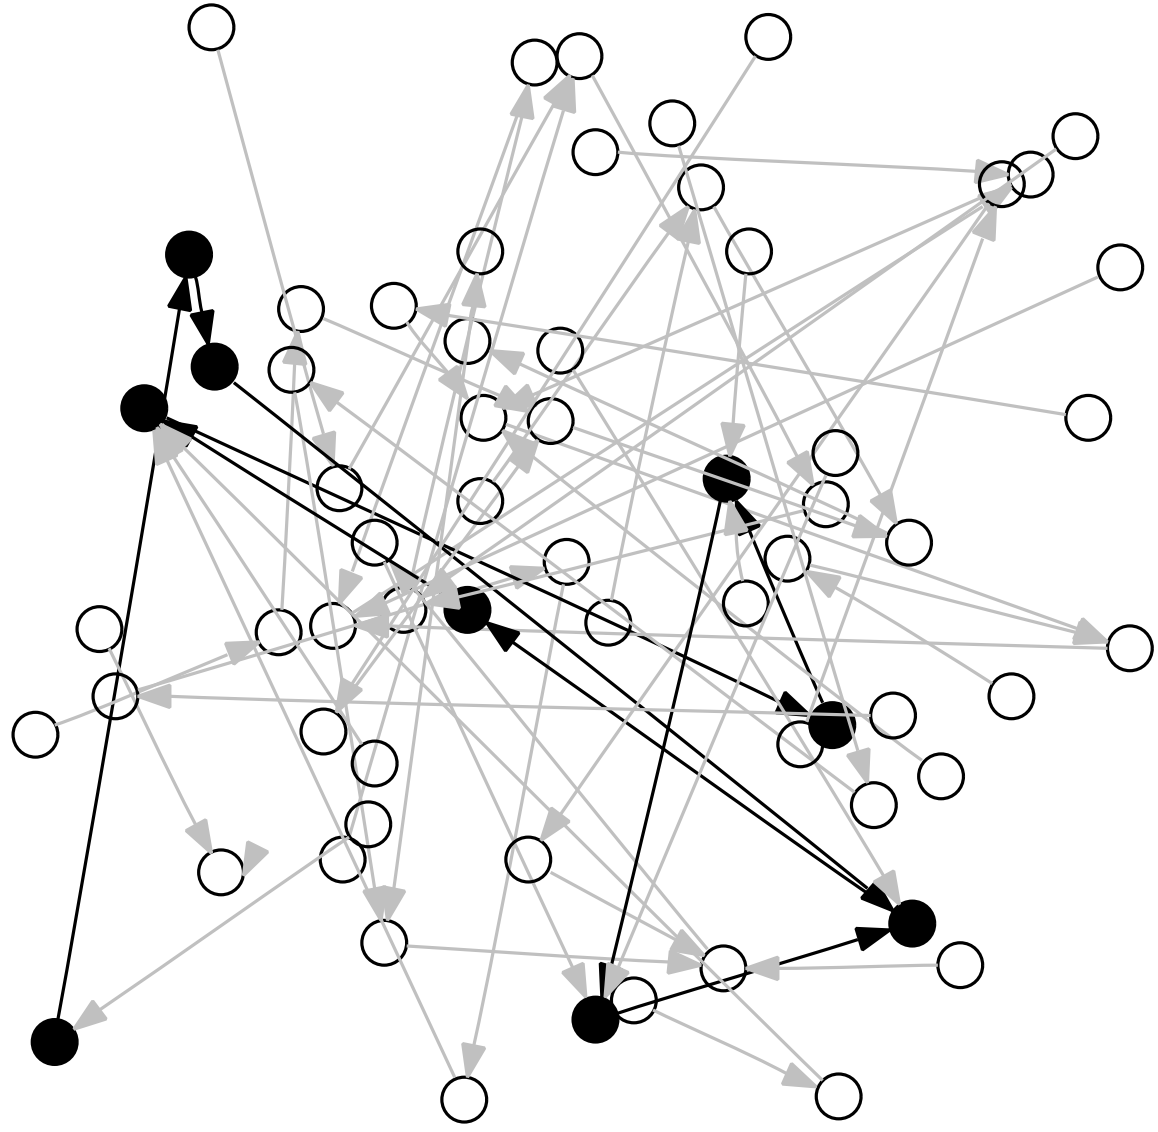

Random walk example

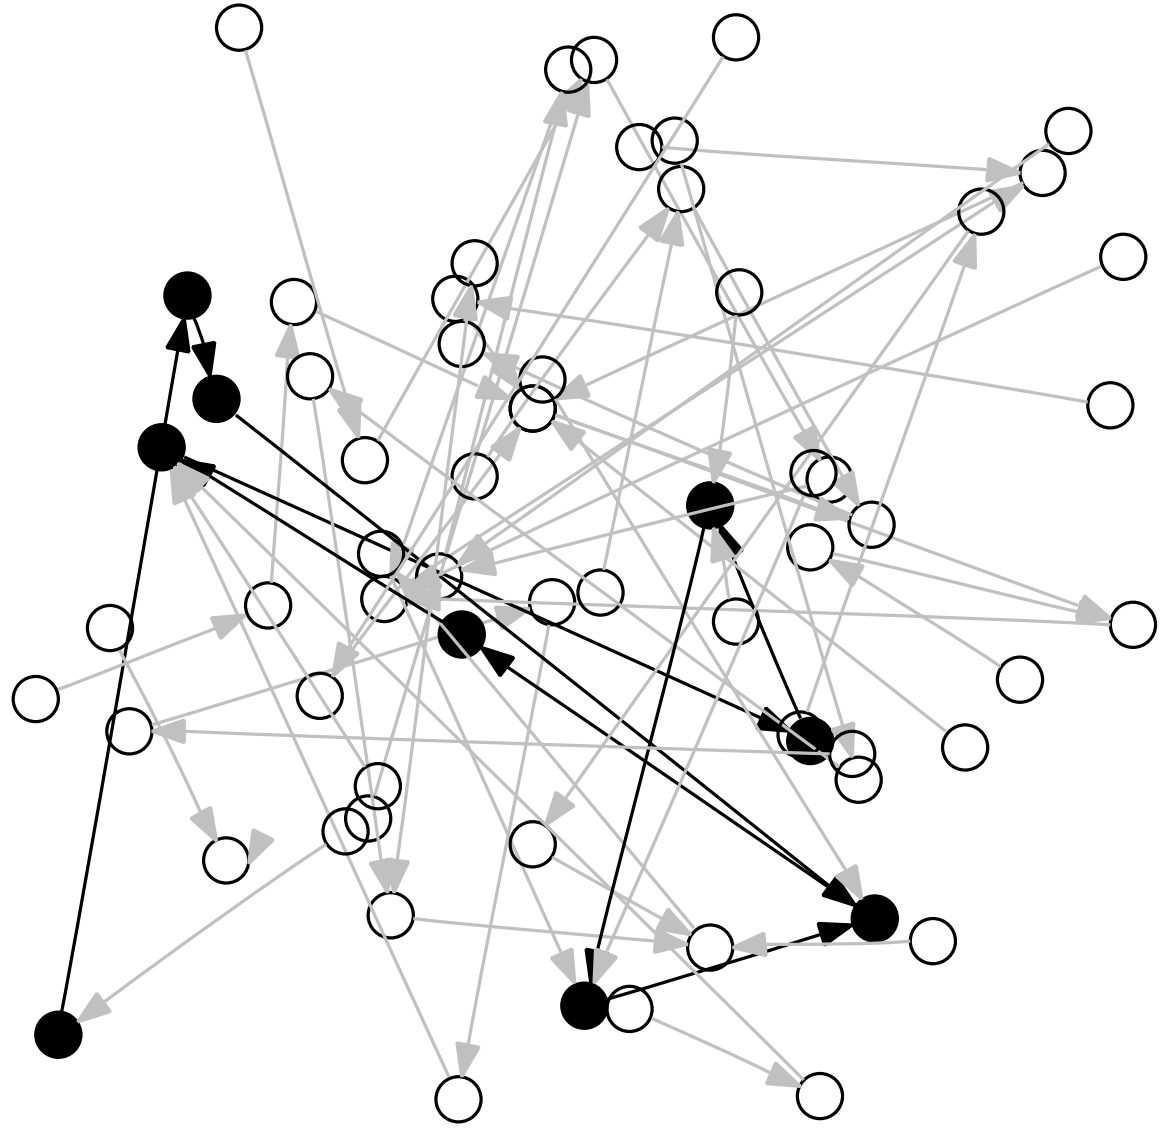

Random walk example

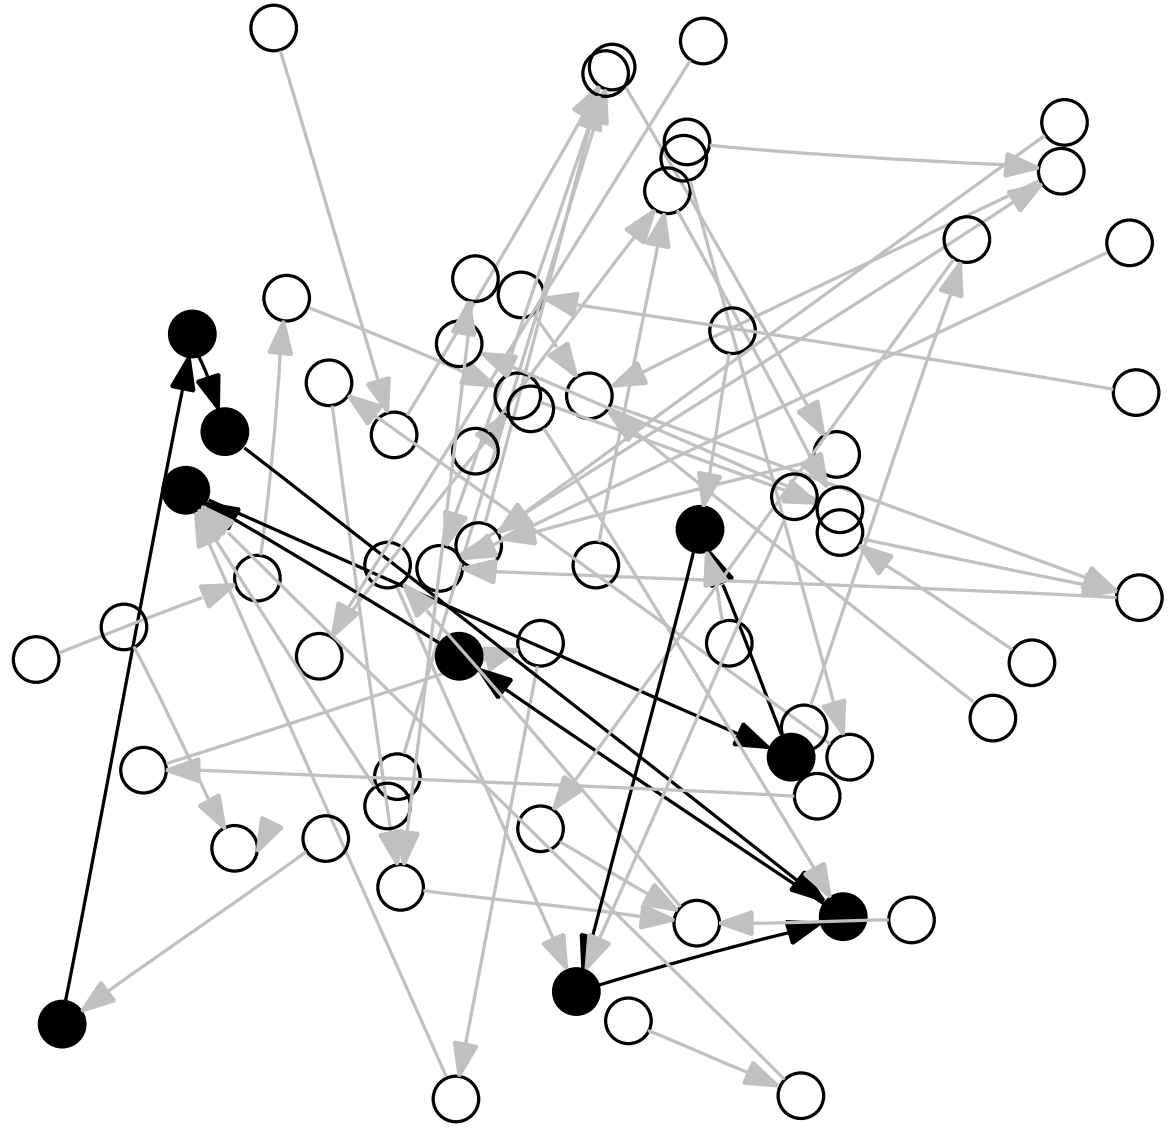

Random walk example

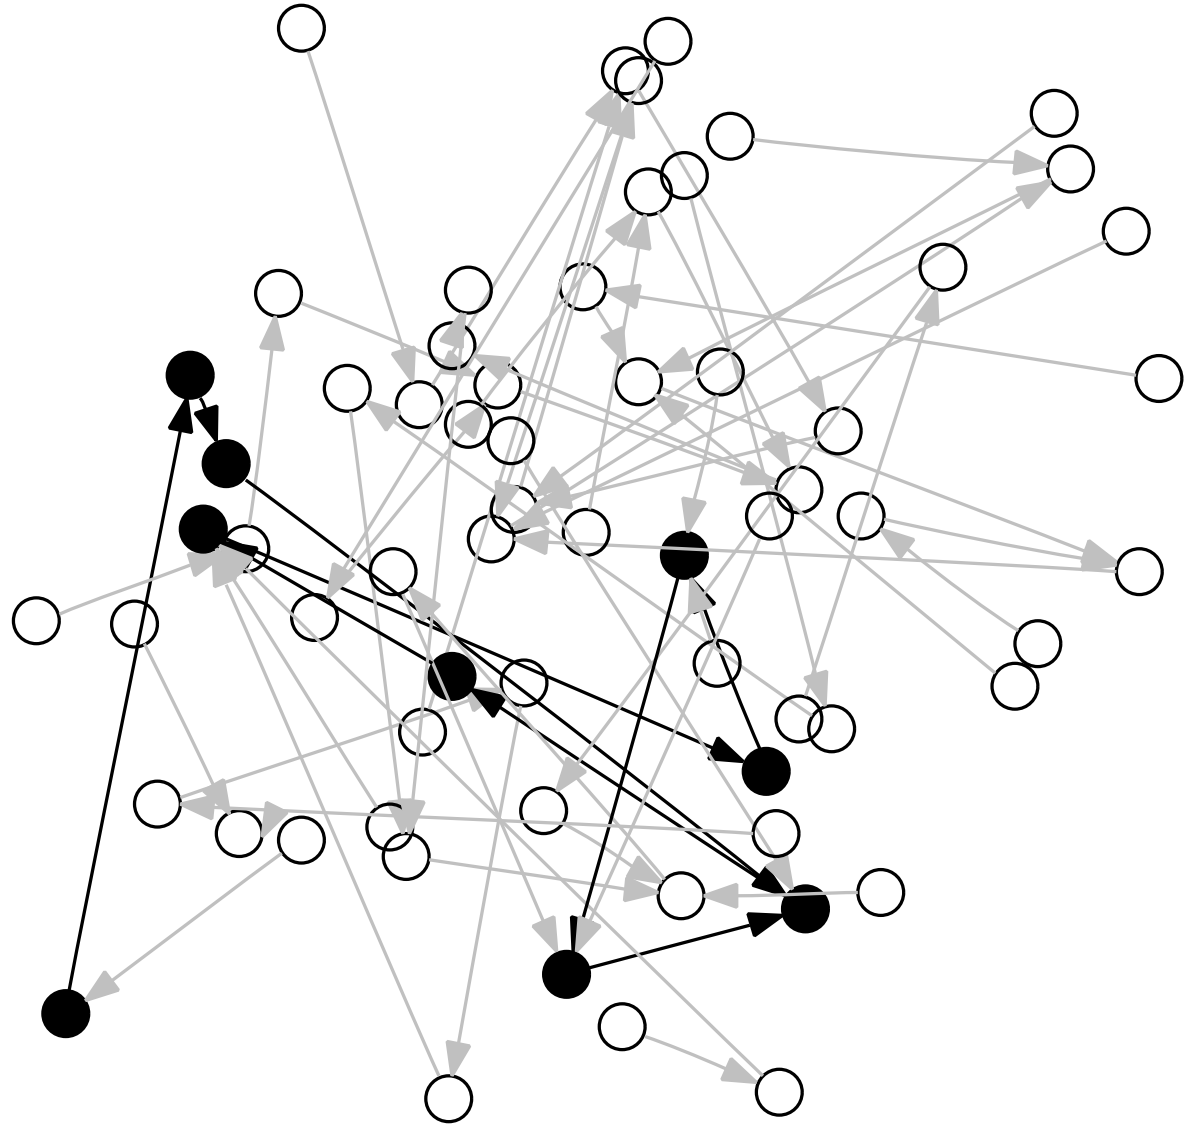

Random walk example

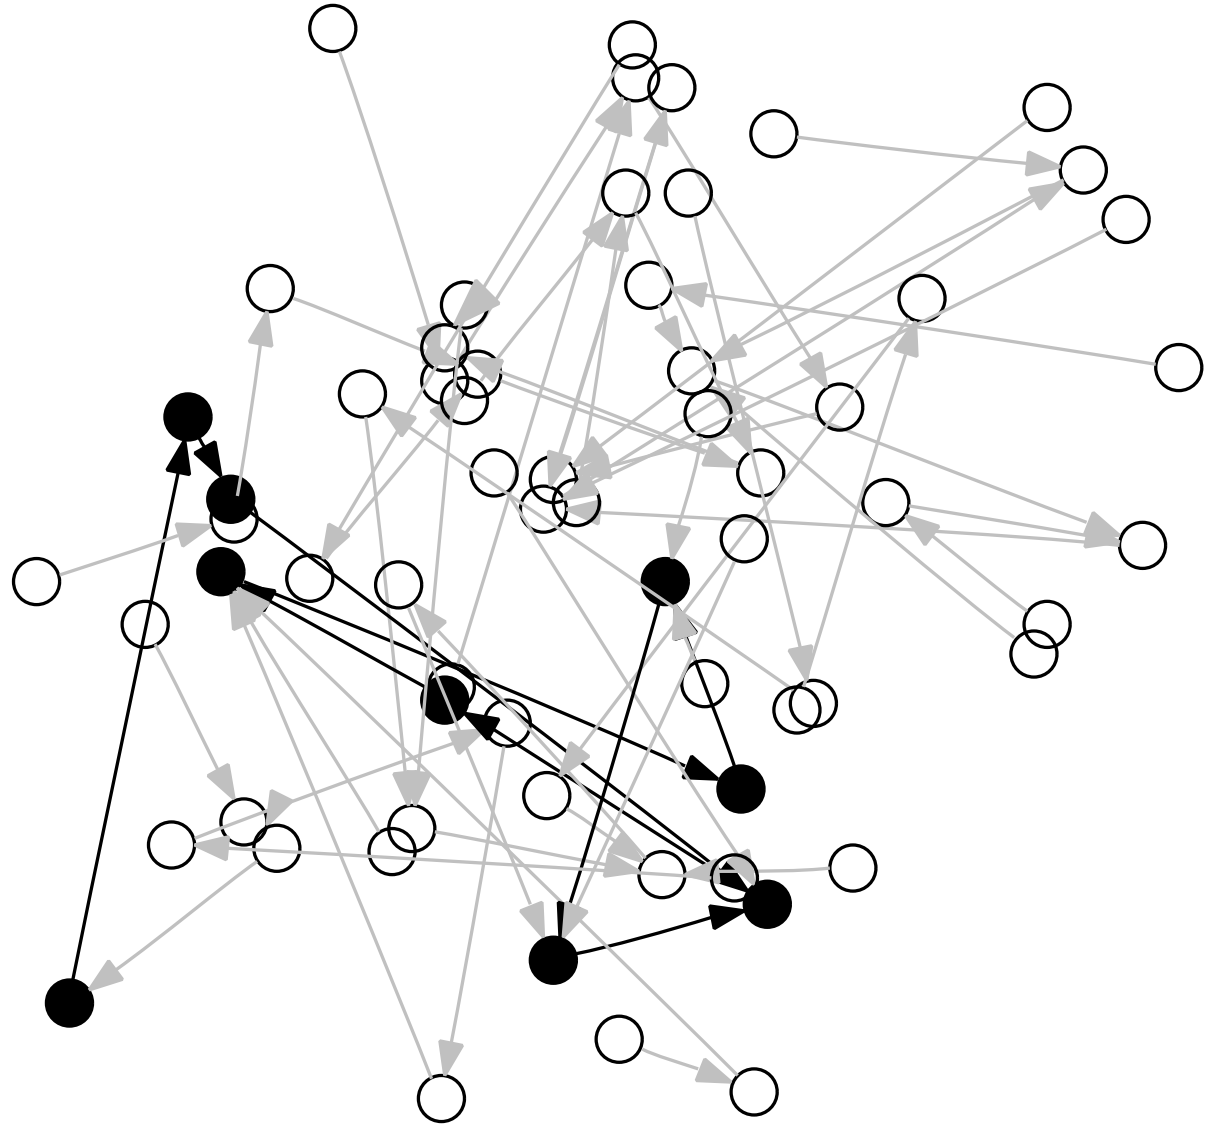

Random walk example

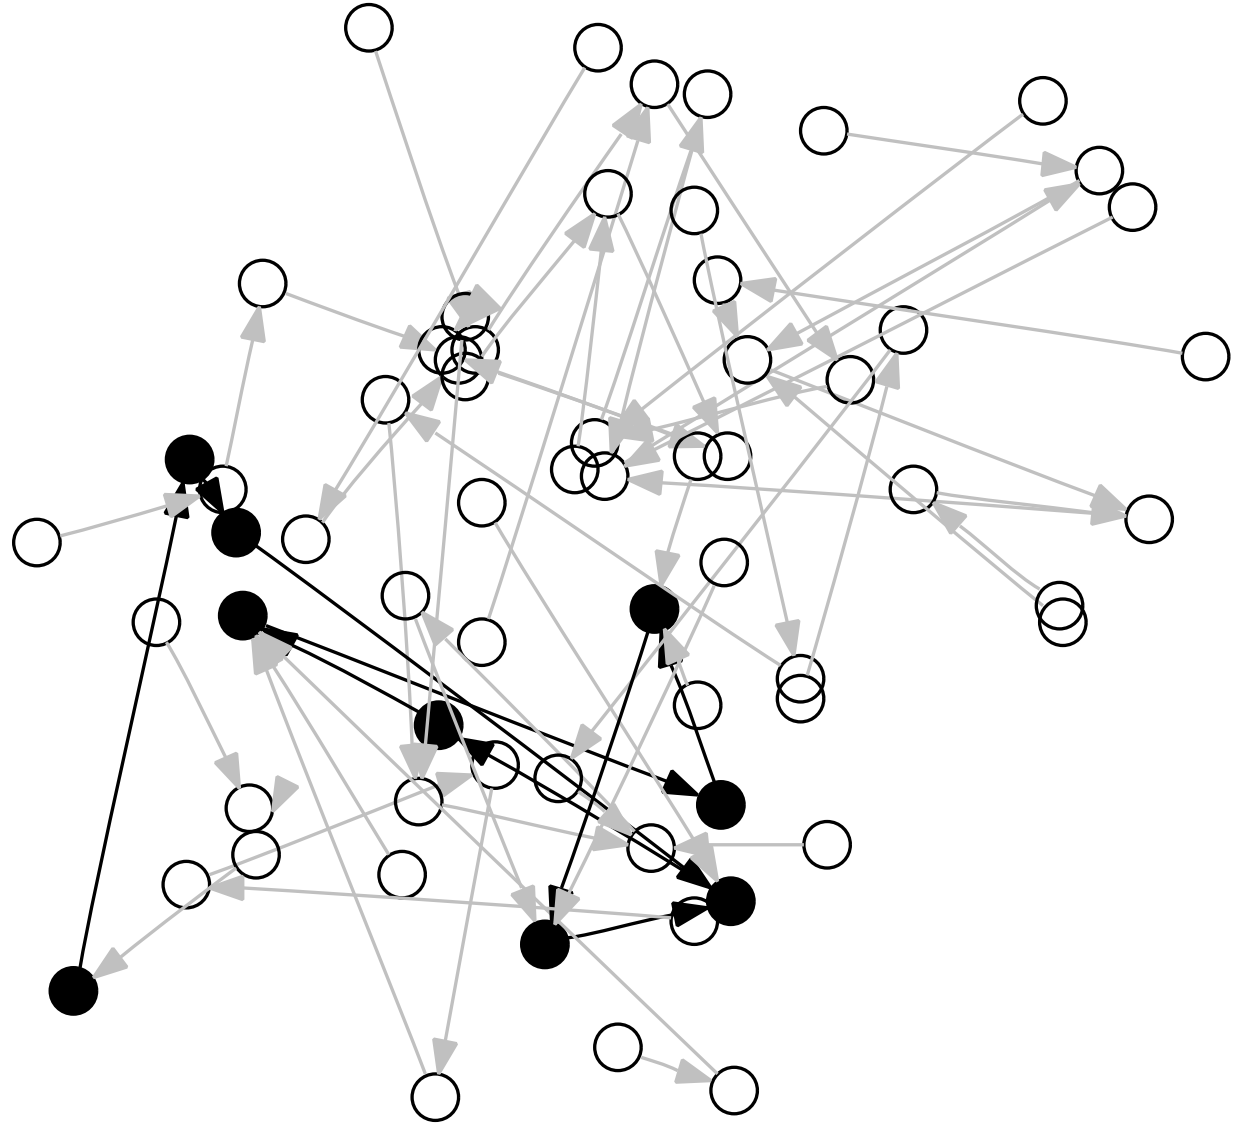

Random walk example

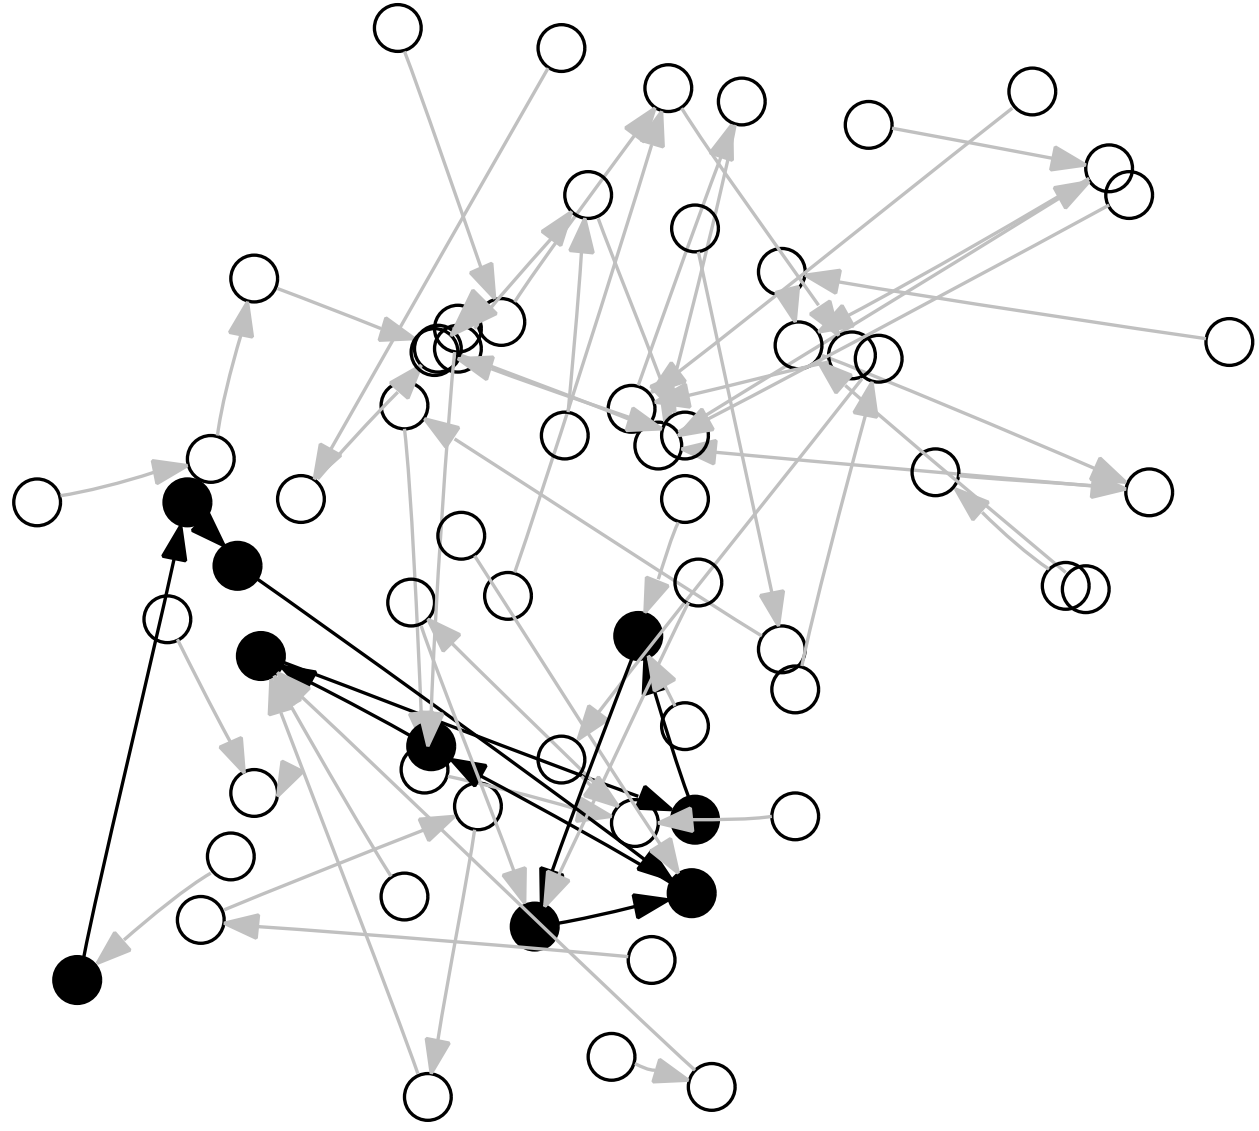

Random walk example

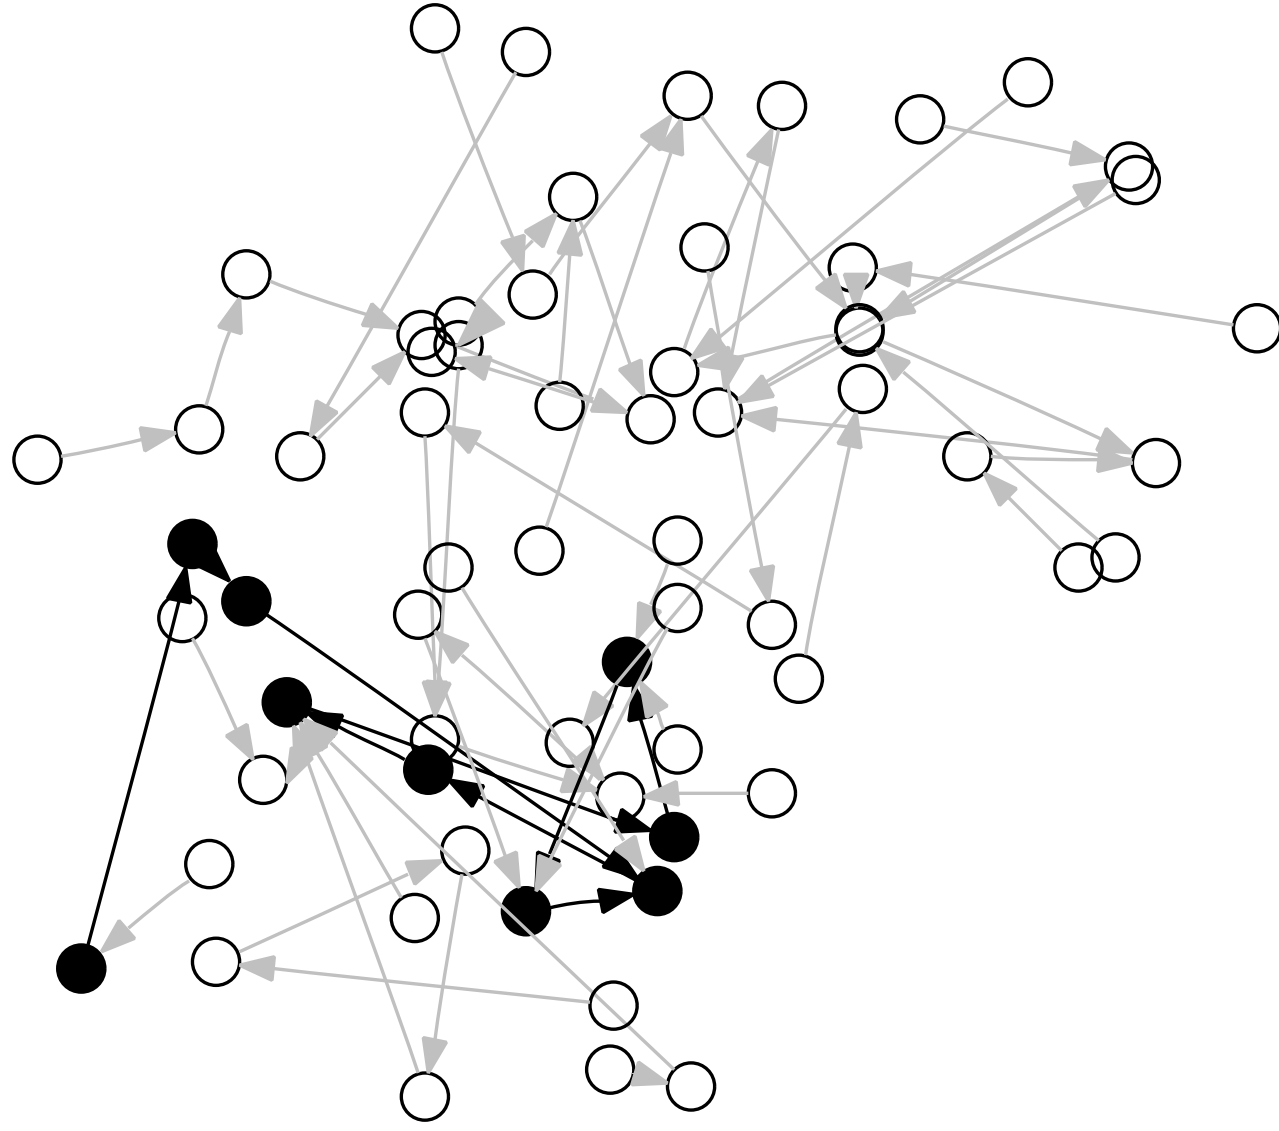

Random walk example

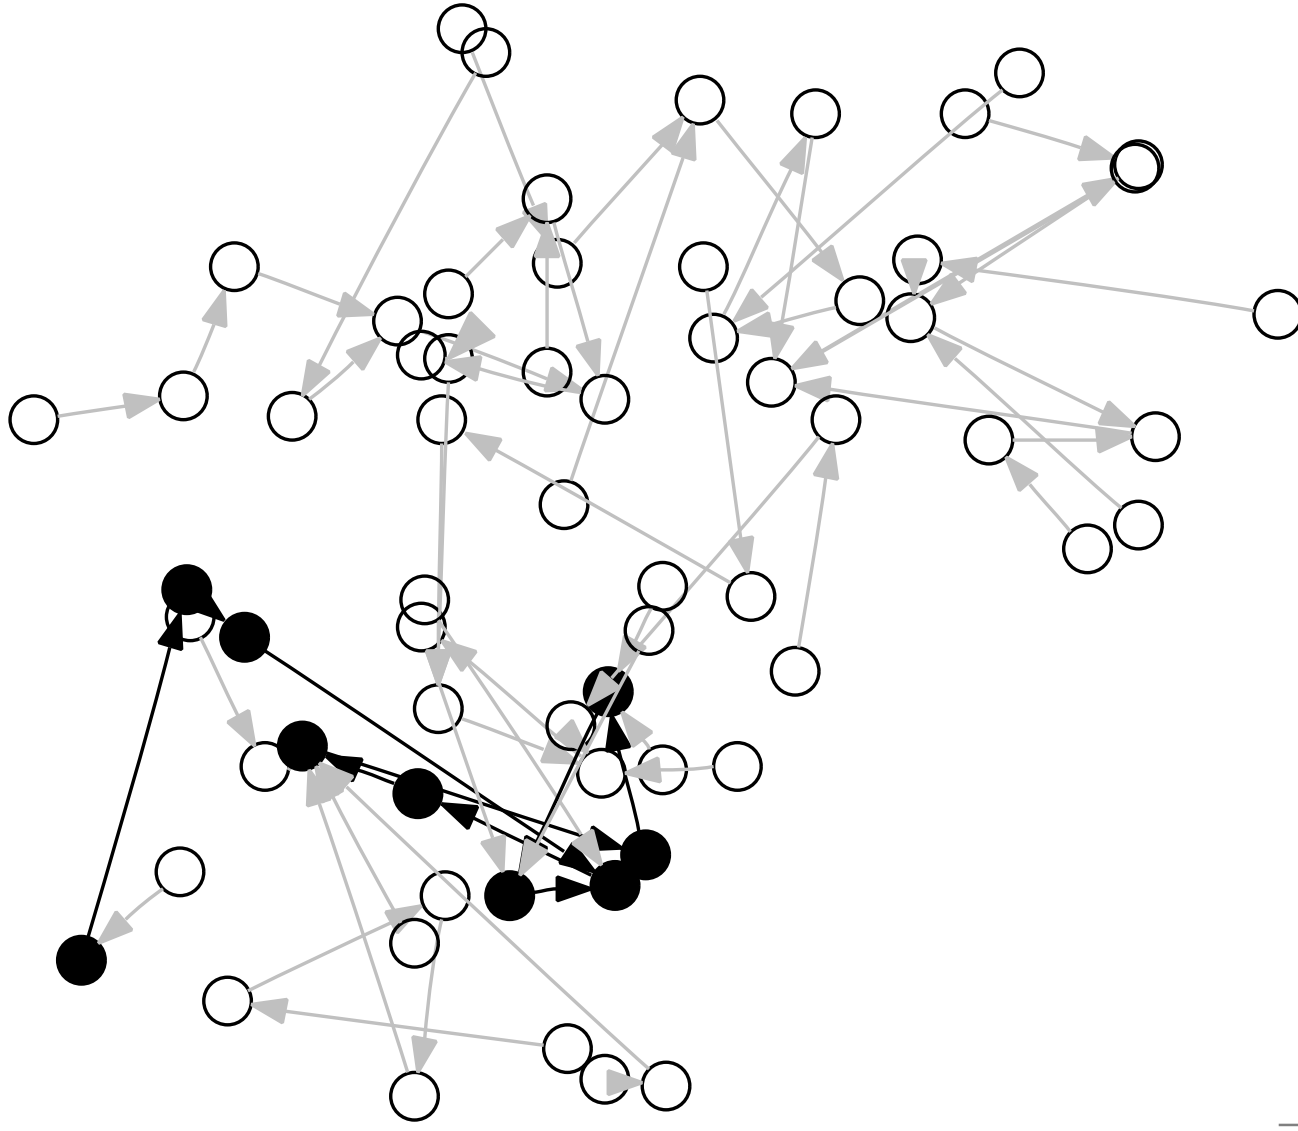

Random walk example

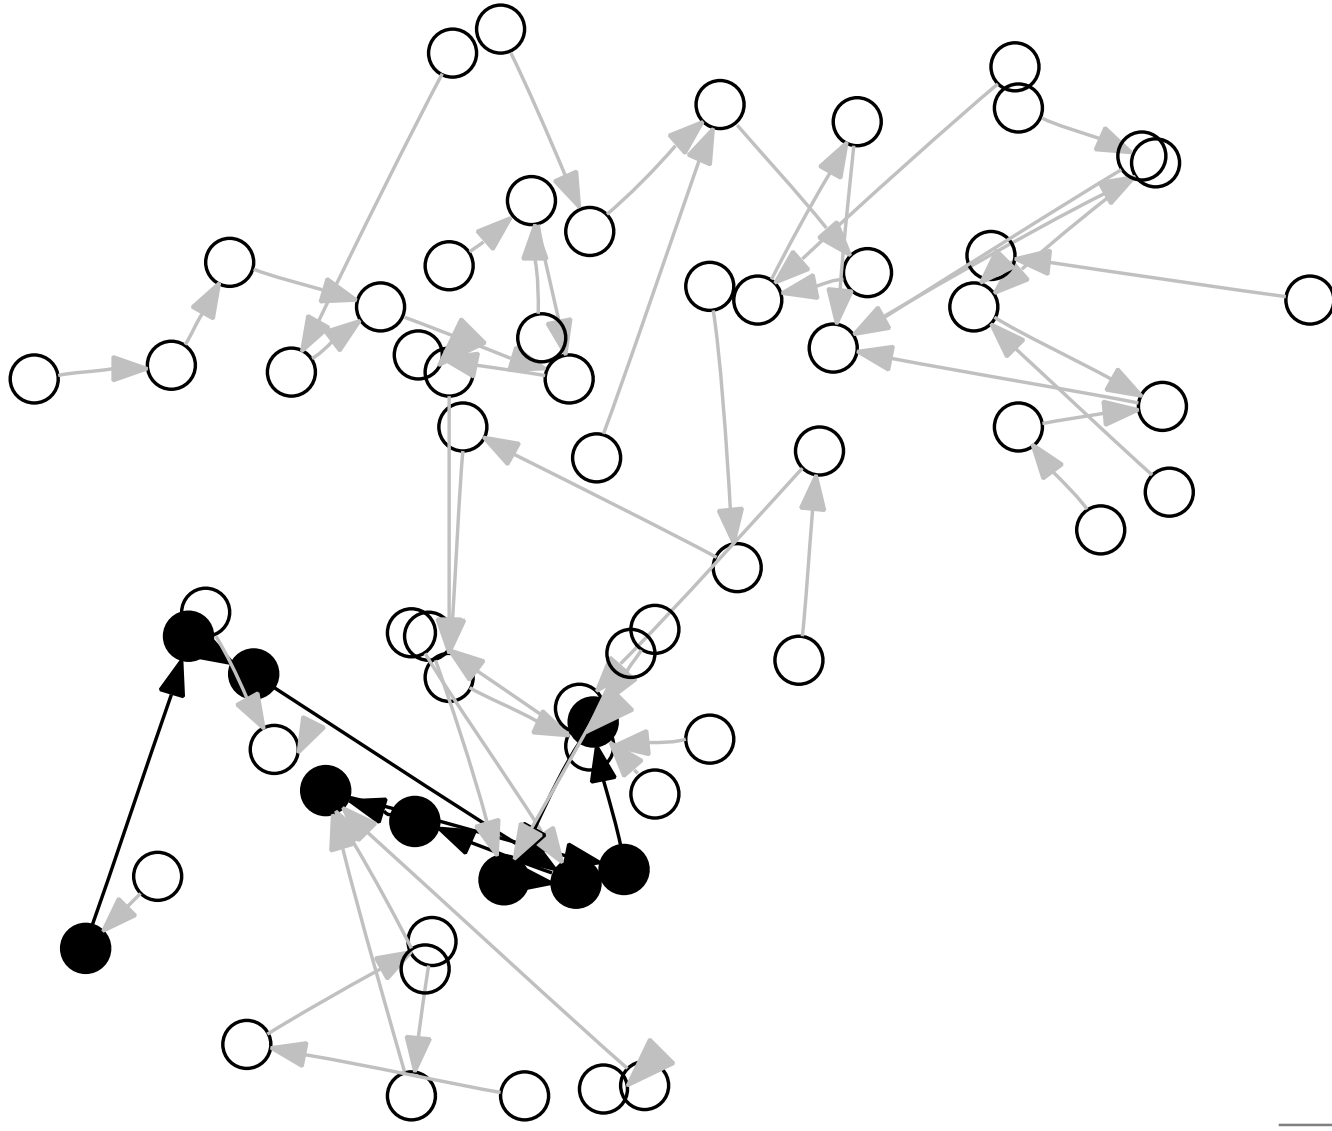

Random walk example

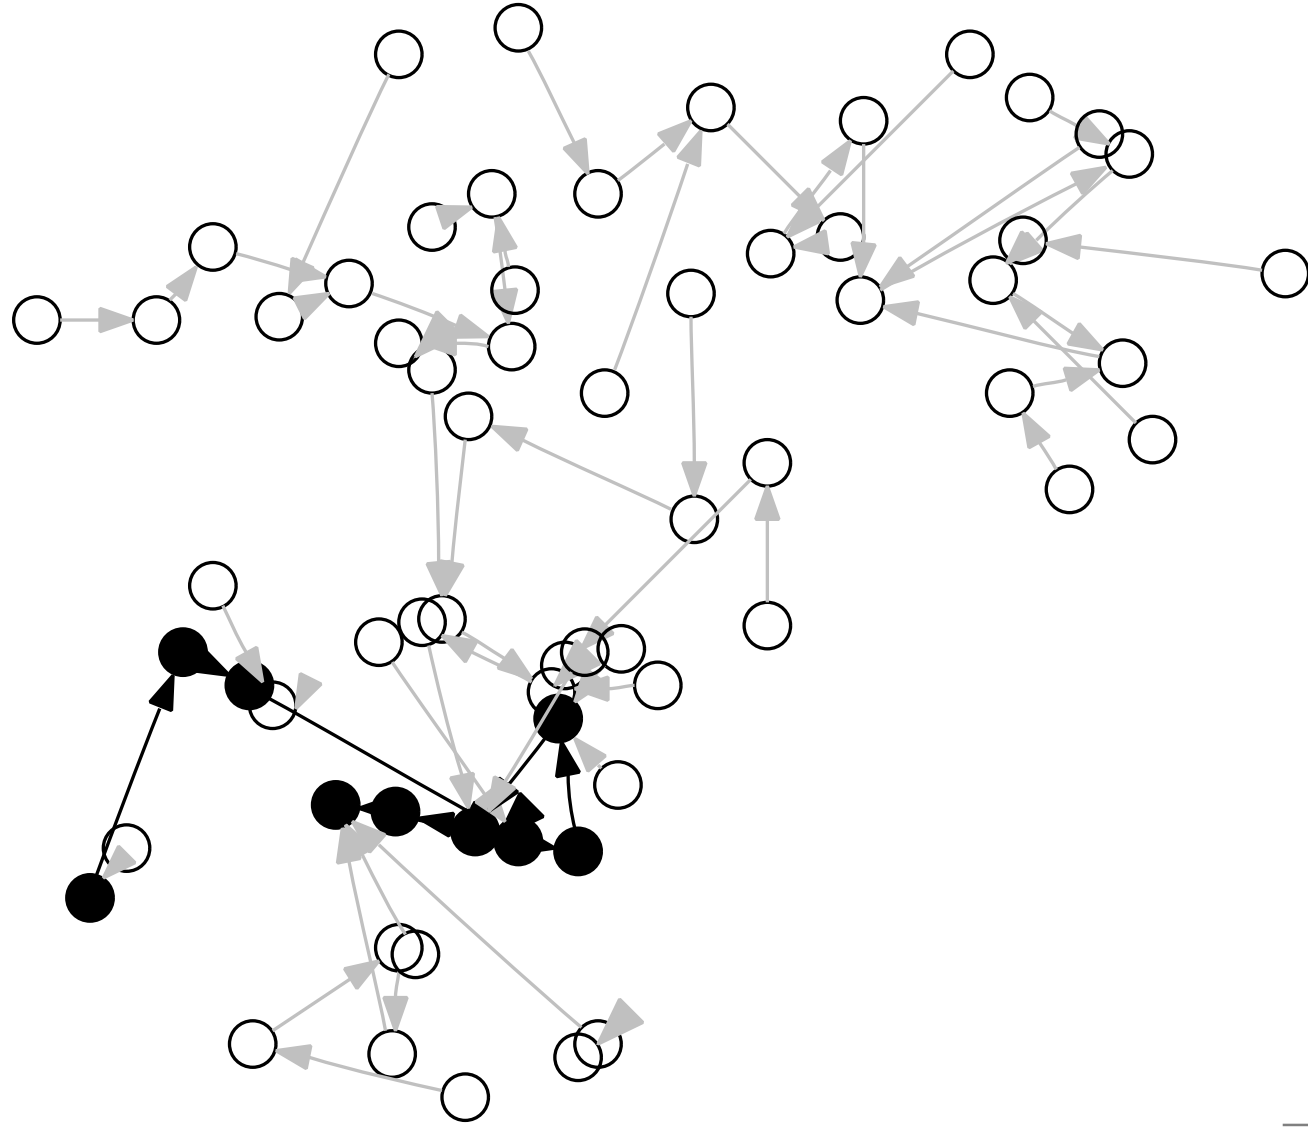

Random walk example

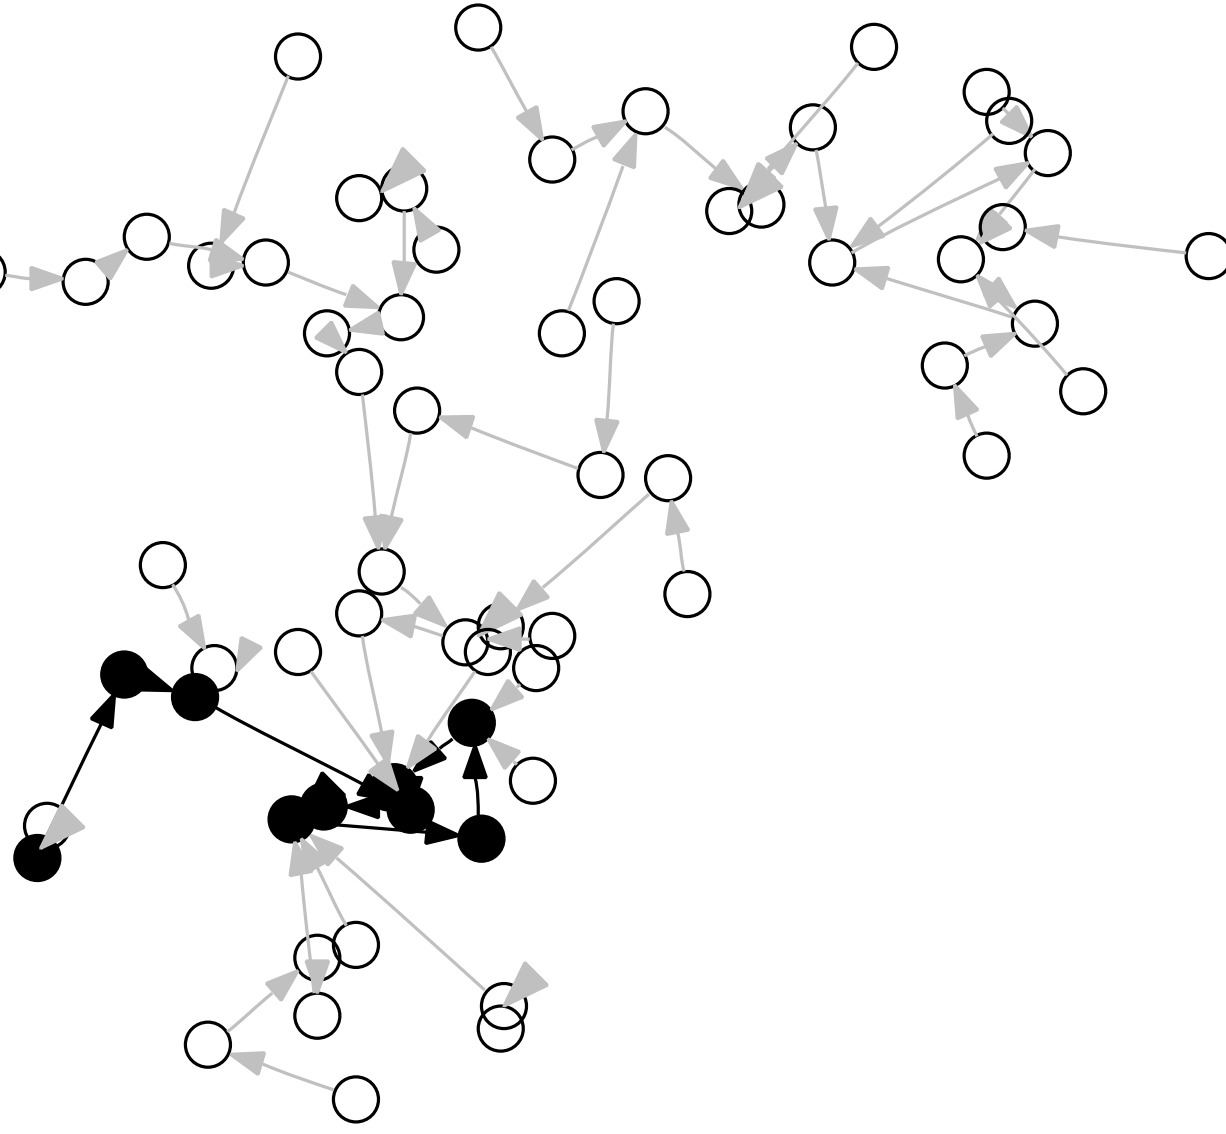

Random walk example

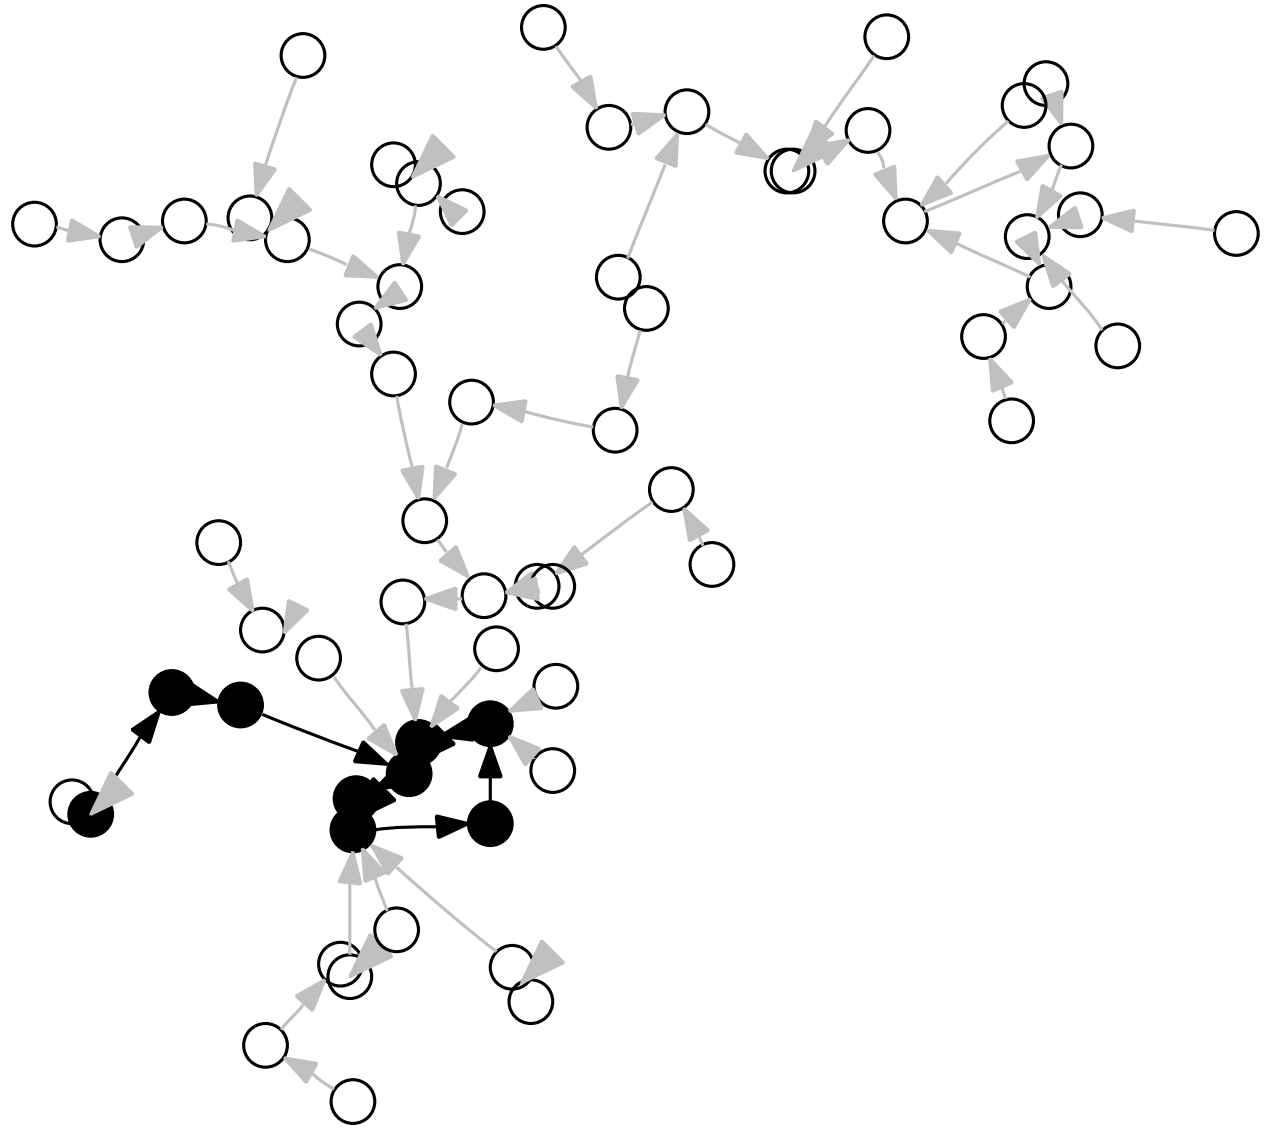

Random walk example

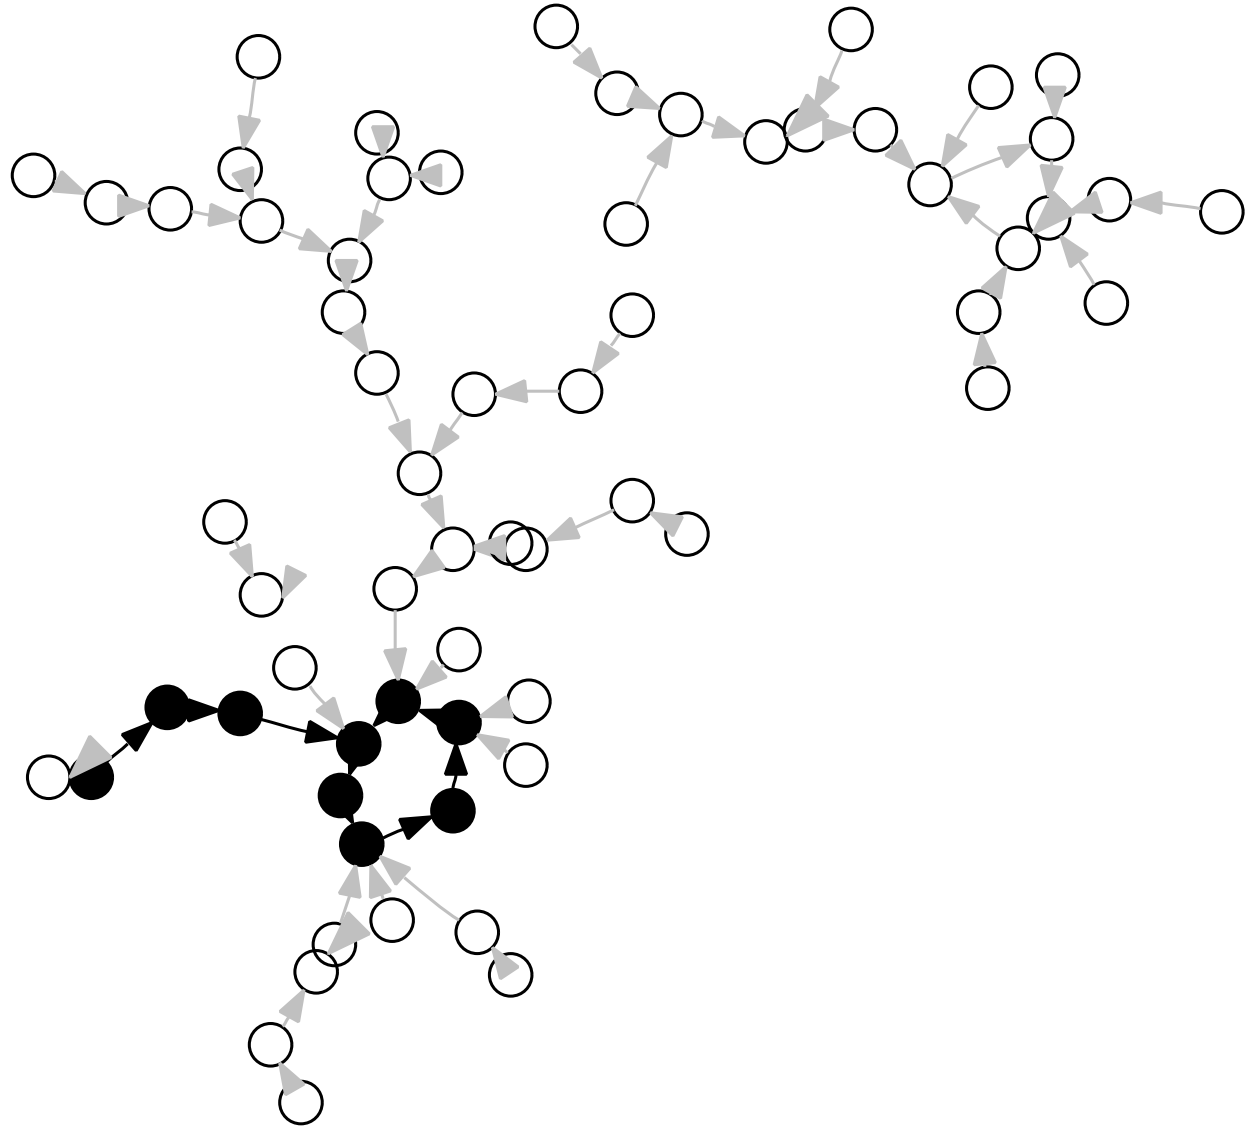

Random walk example

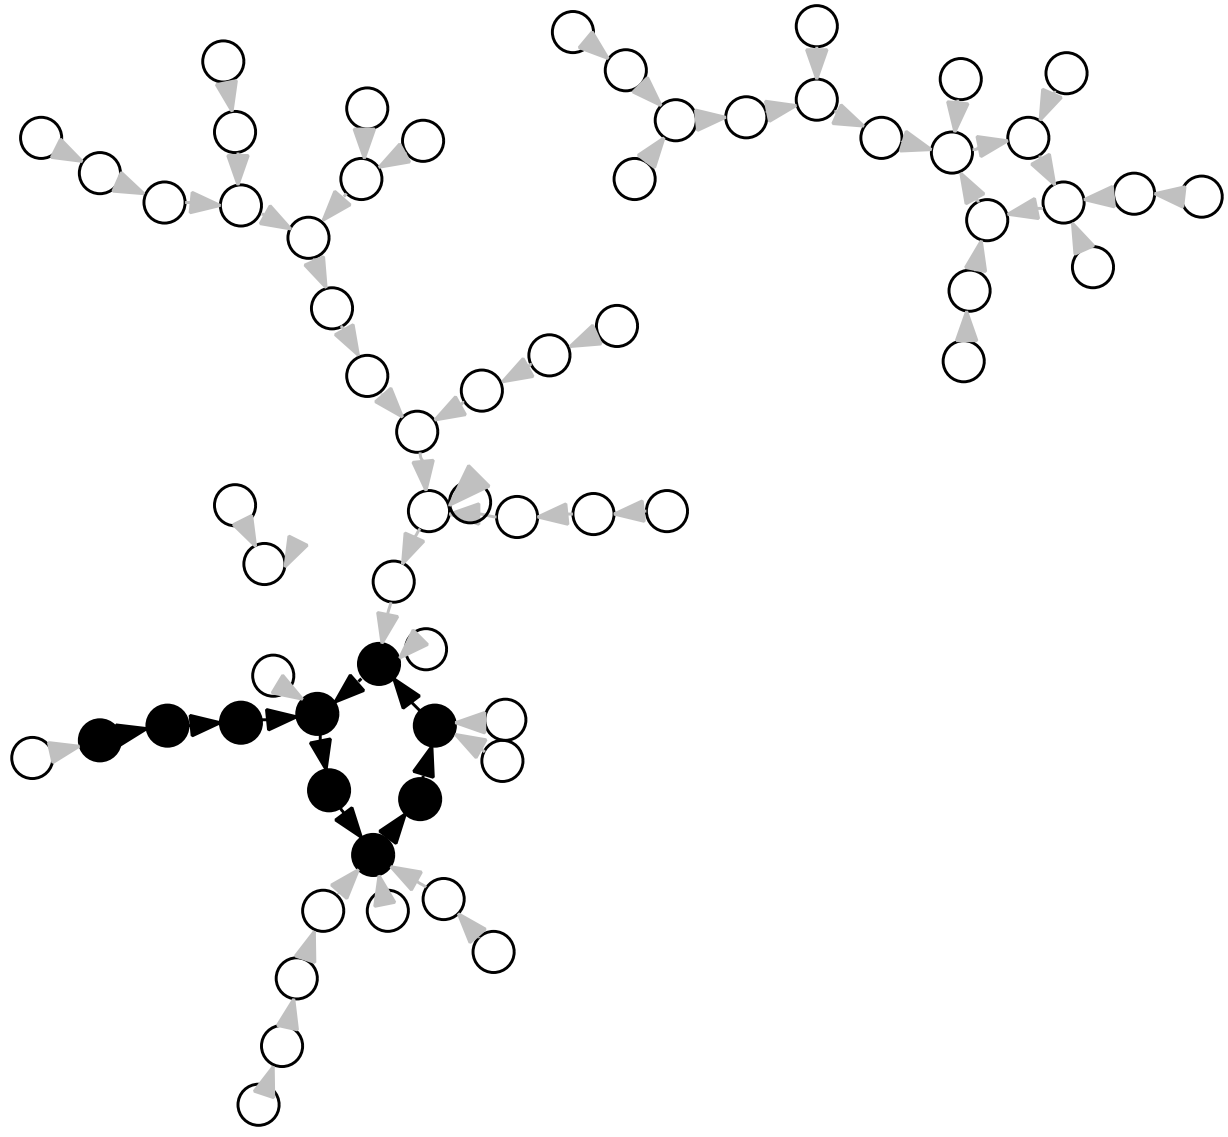

Connection with DLP

- Assume that for each point we know $a_i, b_i \in \mathbb{Z}/\ell\mathbb{Z}$ so that $W_i = [a_i]P + [b_i]Q$.
- Then $W_i = W_j$ means that

$$[a_i]P + [b_i]Q = [a_j]P + [b_j]Q$$

so $[b_i - b_j]Q = [a_j - a_i]P$. If $b_i \neq b_j$ the DLP is solved:

$$a = (a_j - a_i)/(b_i - b_j).$$

- e.g. additive walk: Start with $W_0 = P$ and put

$$f(W_i) = W_i + c_j P + d_j Q,$$

where $j = h(W_i)$.

Parallel rho

- Perform many walks with different starting points but same update function f.
- If two different walks find the same point then their subsequent steps will match.
- Terminate each walk once it hits a distinguished point.
- Attacker chooses frequency and definition of distinguished points. Do not wait for cycle.
- Collect all distinguished points. Two walks ending in same distinguished point solve DLP.

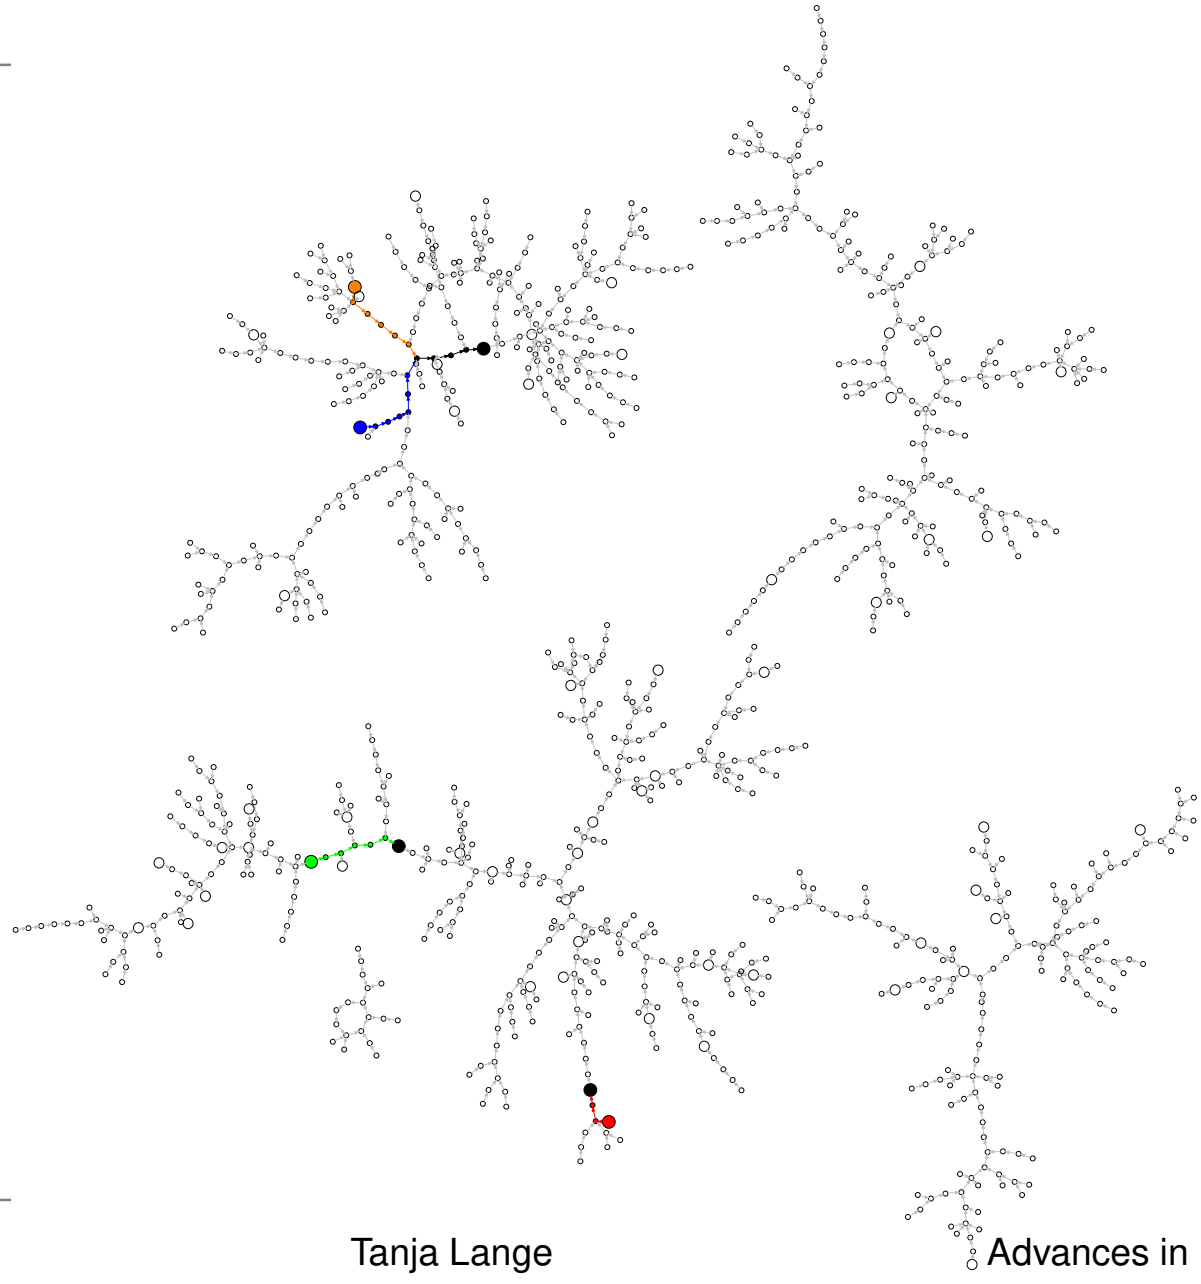

Group law in $E(\mathbb{R})$

$v^2 = u^3 - u$

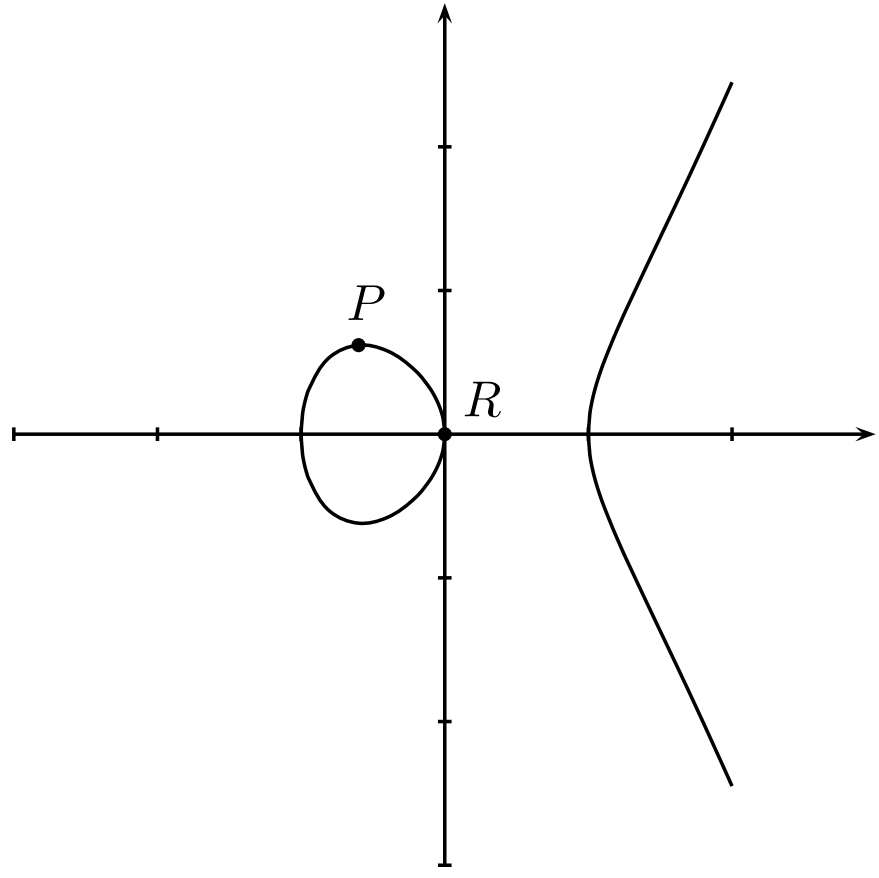

Group law in $E(\mathbb{R})$

$v^2 = u^3 - u$

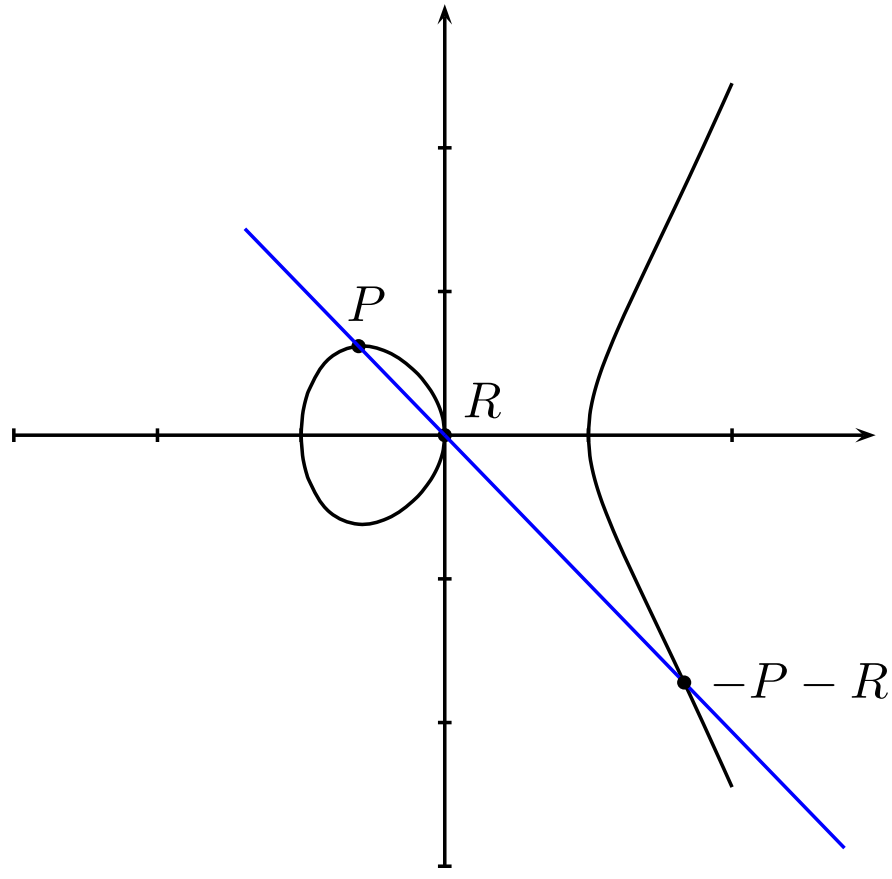

Group law in $E(\mathbb{R})$

$v^2 = u^3 - u$

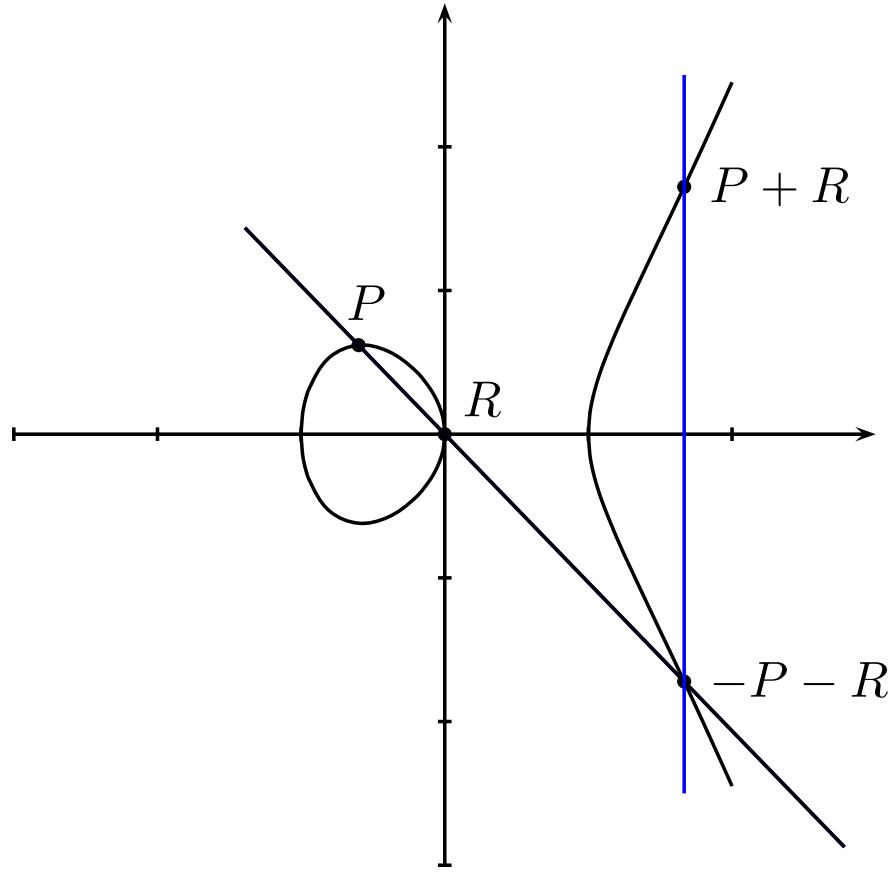

Group law in formulas

$$E : v^2 = u^3 + a_4u + a_6,\ a_i \in \mathbb{F}_q$$

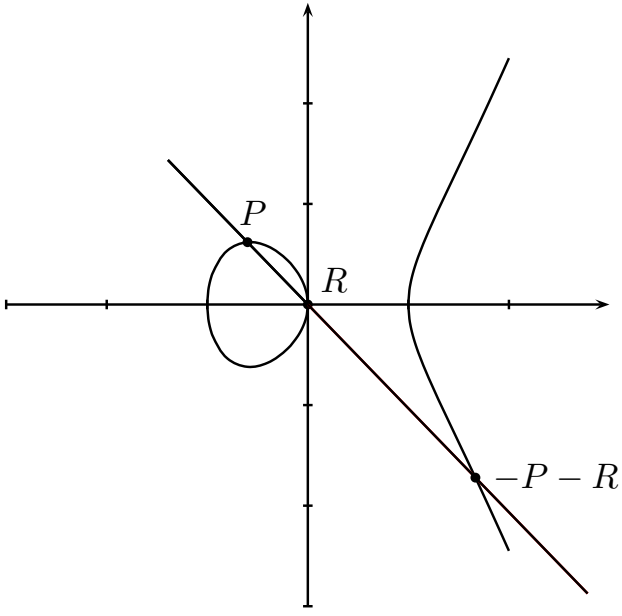

Line $v = \lambda u + \mu$ has slope

$$\lambda = \frac{v_R - v_P}{u_R - u_P}.$$

Equating gives

$$(\lambda u + \mu)^2 = u^3 + a_4u + a_6.$$

This equation has 3 solutions, the u-coordinates of P, R and $-P - R$, thus

$$\begin{aligned}(u - u_P)(u - u_R)(u - u_{-P-R}) &= u^3 - \lambda^2u^2 + (a_4 - 2\lambda\mu)u + a_6 - \mu^2\\ u_{-P-R} &= \lambda^2 - u_P - u_R\end{aligned}$$

Group law in formulas

$$E : v^2 = u^3 + a_4u + a_6,\ a_i \in \mathbb{F}_q$$

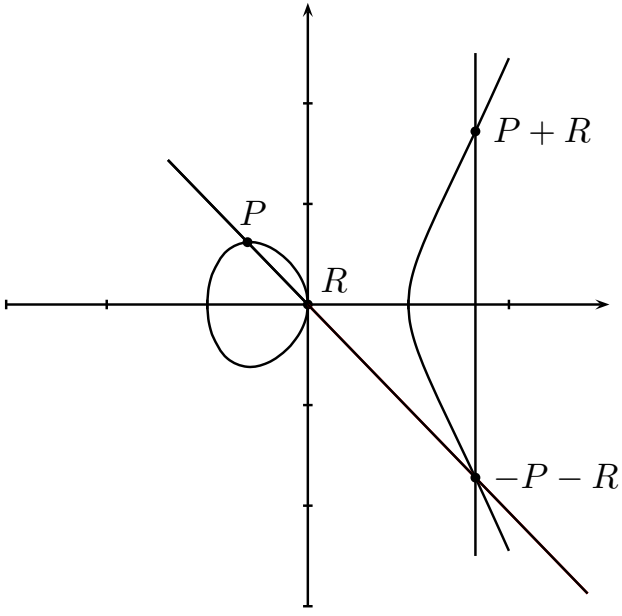

Point P is on line, thus

$$v_P = \lambda u_P + \mu, \text{ i.e.}$$
$$\mu = v_P - \lambda u_P,$$

and

$$\begin{aligned} v_{-P-R} &= \lambda u_{-P-R} + \mu \\ &= \lambda u_{-P-R} + v_P - \lambda u_P \\ &= \lambda(u_{-P-R} - u_P) + v_P \end{aligned}$$

Point $P+R$ has the same u-coordinate but negative v-coordinate:

$$u_{P+R} = \lambda^2 - u_P - u_R, \quad v_{P+R} = \lambda(u_P - u_{P+R}) - v_P$$

Group law in formulas

$$E : v^2 = u^3 + a_4u + a_6,\ a_i \in \mathbb{F}_q$$

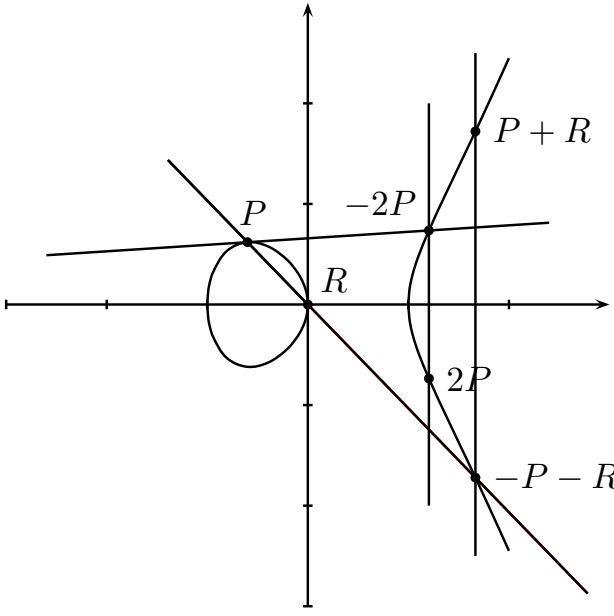

In general, for $(u_P, v_P) \neq (u_R, -v_R)$:

$$\begin{aligned}(u_P, v_P) + (u_R, v_R) &= \\ &= (u_{P+R}, v_{P+R}) = \\ &= (\lambda^2 - u_P - u_R, \lambda(u_P - u_{P+R}) - v_P),\end{aligned}$$

where

$$\lambda = \begin{cases} (v_R - v_P)/(u_R - u_P) & \text{if } u_P \neq u_R, \\ (3u_P^2 + a_4)/(2v_P) & \text{else.} \end{cases}$$

$\Rightarrow$ Addition and Doubling need

1 I, 2M, 1S and 1 I, 2M, 2S, respectively

Negation and rho

- $W = (x, y)$ and $-W = (x, -y)$ have same x-coordinate. Search for x-coordinate collision.
- Search space for collisions is only $\lceil \ell/2 \rceil$; this gives factor $\sqrt{2}$ speedup ... if $f(W_i) = f(-W_i)$.
- To ensure $f(W_i) = f(-W_i)$ define $j = h(|W_i|)$ and $f(W_i) = |W_i| + c_j P + d_j Q$.
  Define $|W_i|$ as, e.g., lexicographic minimum of $W_i, -W_i$.

Negation and rho

- $W = (x, y)$ and $-W = (x, -y)$ have same x-coordinate. Search for x-coordinate collision.
- Search space for collisions is only $\lceil \ell/2 \rceil$; this gives factor $\sqrt{2}$ speedup ... if $f(W_i) = f(-W_i)$.
- To ensure $f(W_i) = f(-W_i)$ define $j = h(|W_i|)$ and $f(W_i) = |W_i| + c_j P + d_j Q$.
  Define $|W_i|$ as, e.g., lexicographic minimum of $W_i, -W_i$.
- Problem: this walk can run into fruitless cycles!
- E.g. if $|W_{i+1}| = -W_{i+1}$ and $h(|W_{i+1}|) = j = h(|W_i|)$ then

$$\begin{aligned} W_{i+2} &= f(W_{i+1}) = -W_{i+1} + c_j P + d_j Q \\ &= -(|W_i| + c_j P + d_j Q) + c_j P + d_j Q = -|W_i|, \end{aligned}$$

so $|W_{i+2}| = |W_i|$, so $W_{i+3} = W_{i+1}$, so $W_{i+4} = W_{i+2}$ etc.

Fruitless cycles

- If h maps to r different values then expect fruitless cycles of length 2 to occur with probability $1/(2r)$ at each step.
- Longer cycles occur less frequently.
- Current ECDL record, achieved 07.2009 by Bos–Kaihara–Kleinjung–Lenstra–Montgomery, "PlayStation 3 computing breaks 2^{60} barrier – 112-bit prime ECDLP solved" breaks standard curve over $\mathbb{F}_p$, where $p = (2^{128} - 3)/(11 \cdot 6949)$.
- "We did not use the common negation map since it requires branching and results in code that runs slower in a SIMD environment." (All modern CPUs are SIMD.)
- Several more papers justifying not using negation map.

Correct use of the negation map I

- Our software (Bernstein-L-Schwabe, PKC 2011) solves random ECDL on the same curve (with no precomputation) in 35.6 PS3 years on average.
- For comparison, software from record computation uses 65 PS3 years on average.
- Several levels of speedups, starting with fast arithmetic mod p and continuing up through rho.
- Most important speedup: We use the negation map.
- Extra cost in each iteration:
  - extract bit of "s"(normalized y, needed anyway);
  - expand bit into mask;
  - use mask to conditionally replace (s, y) by $(-s, -y)$.
  - $\approx 1.5\%$ of total; no conditional branches.

Problems with negation map

- Bos–Kleinjung–Lenstra say that “on average more elliptic curve group operations are required per step of each walk. This is unavoidable” etc.
- Specifically: If the precomputed additive-walk table has r points, need 1 extra doubling to escape a cycle after $\approx 2r$ additions. And more: “cycle reduction” etc.
- Bos–Kleinjung–Lenstra say that the benefit of large r is “wiped out by cache inefficiencies.”

Correct use of the negation map II

- There's really no problem here!
- We use $r = 2048$. $1/(2r) = 1/4096$; negligible.
- Recall: p has 112 bits. 28 bytes for table entry (x, y).
- We expand to 36 bytes to accelerate arithmetic. We compress to 32 bytes by insisting on small x, y; very fast initial computation.
- Only 64KB for table. Our Cell table-load cost: 0, overlapping loads with arithmetic. No "cache inefficiencies."

What about fruitless cycles?

- We run 45 iterations. We then save s; run 2 slightly slower iterations tracking minimum (s, x, y); then double tracked (x, y) if new s equals saved s.
- Occasionally replace 2 by 12 to detect 4-cycles, 6-cycles. Such cycles almost too rare to worry about, but detecting them has completely negligible cost.
- Maybe fruitless cycles waste some of the 47 iterations. . . . but this is infrequent. Lose $\approx$ 0.6% of all iterations.
- Tracking minimum isn't free, but most iterations skip it! Same for final s comparison. Still no conditional branches. Overall cost $\approx 1.3\%$.
- Doubling occurs for only $\approx 1/4096$ of all iterations. We use SIMD quite lazily here; overall cost $\approx 0.6\%$. Can reduce this cost further.

Experiments

- To confirm iteration effectiveness we have run many experiments on $y^2 = x^3 - 3x + 9$ over the same $\mathbb{F}_p$, using smaller-order P.
- Results matched DL cost predictions.
- Conclusions of this part:
  Sensible use of negation, with or without SIMD, has negligible impact on cost of each iteration. Impact on number of iterations is almost exactly $\sqrt{2}$. Overall benefit is extremely close to $\sqrt{2}$.
- Lessons
  - Breaking the ECDLP is no harder than textbooks tell you – but that's pretty hard.
  - If a system gets broken it's not the fault of the ECDLP.

How are ECDSA systems broken?

- Anybody can produce an $R = kP$. Alice's private key is only used in $s \equiv k^{-1}(h(m) + r \cdot a) \bmod \ell$.
- Can fake signatures if one can break the DLP, i.e., if one can compute a from P_A. . . . sometimes attacks are easier. . .

How are ECDSA systems broken?

- Anybody can produce an $R = kP$. Alice's private key is only used in $s \equiv k^{-1}(h(m) + r \cdot a) \bmod \ell$.
- Can fake signatures if one can break the DLP, i.e., if one can compute a from P_A. . . . sometimes attacks are easier. . .
- If k is known for some $m, (r, s)$ then $a \equiv (sk - h(m))/r \bmod \ell$.

How are ECDSA systems broken?

- Anybody can produce an $R = kP$. Alice's private key is only used in $s \equiv k^{-1}(h(m) + r \cdot a) \bmod \ell$.
- Can fake signatures if one can break the DLP, i.e., if one can compute a from P_A. . . . sometimes attacks are easier. . .
- If k is known for some $m, (r, s)$ then $a \equiv (sk - h(m))/r \bmod \ell$.
- If two signatures $m_1, (r, s_1)$ and $m_2, (r, s_2)$ have the same value for r: assume $k_1 = k_2$; observe $s_1 - s_2 = k_1^{-1}(h(m_1) + ra - (h(m_2) + ra))$; compute $k = (s_1 - s_2)/(h(m_1) - h(m_2))$. Continue as above.

How are ECDSA systems broken?

- Anybody can produce an $R = kP$. Alice's private key is only used in $s \equiv k^{-1}(h(m) + r \cdot a) \bmod \ell$.
- Can fake signatures if one can break the DLP, i.e., if one can compute a from P_A. . . . sometimes attacks are easier. . .
- If k is known for some $m, (r, s)$ then $a \equiv (sk - h(m))/r \bmod \ell$.
- If two signatures $m_1, (r, s_1)$ and $m_2, (r, s_2)$ have the same value for r: assume $k_1 = k_2$; observe $s_1 - s_2 = k_1^{-1}(h(m_1) + ra - (h(m_2) + ra))$; compute $k = (s_1 - s_2)/(h(m_1) - h(m_2))$. Continue as above.
- If bits of many k's are known (biased PRNG) can attack $s \equiv k^{-1}(h(m) + r \cdot a) \bmod \ell$ as hidden number problem using lattice basis reduction.

Avoid trouble and be faster

aka “High-speed high-security signatures” (Bernstein–Duif–L–Schwabe–Yang, CHES 2011).

- Fast signature verification. Intel Nehalem/Westmere CPUs: 280880 cycles for verification. (official eBATS measurement.)
- Even faster batch verification: < 134000 cycles/signature to verify 64 signatures of 64 messages under 64 public keys. 4×2.4GHz: 71000 verif/second!
- Very fast signing. 88328 cycles for signing. (eBATS measurement.) 4×2.4GHz: 108000 signs/second!
- Fast key generation. Almost as fast as signing.
- High security level > 2^{128} (without quantum computers).
- Small signatures and keys.

Bonus features

- No secret array indices. No information flow from secret data to addresses. $\Rightarrow$ No cache-timing attacks.
- No secret branch conditions. No information flow from secret data to branch unit.
- Collision resilience. Hash collisions do not break this signature system.
- Foolproof session keys. Signing is deterministic.
- Paper contains new representation of points which is faster for adding precomputed points.
- Curve is an Edwards curve, so it looks nicer :–)
- New efficient point compression and decompression for Edwards curves.

Some details

- Secret key of user produces a used in $P_A = aP$ and bitstring b.
- Nonce k is computed as $k =$hash(b, m).
- $R = kP$ as before but h is made dependent on R and P_A as $h =$hash(R, P_A, m).
- Signing and verification avoid inversions $s \equiv k + ha \bmod \ell$.
- Verification: $sP \stackrel{?}{=} R + hP_A$.
- Batch verification of (R_i, s_i) with $h_i =$hash(R_i, P_{A_i}, m_i) picks random scalars z_i and checks whether

$$(-\sum_i z_i s_i \bmod \ell)P + \sum_i z_i R_i + \sum_i (z_i h_i \bmod \ell)P_{A_i} = 0.$$

Thank you for your attention!

Addition on Edwards curves

Let K be a field with $2 \neq 0$. Let $d \in K$ with $d \neq 0, 1$. y Edwards curve (nice form of elliptic curve):

$$\{(x, y) \in K \times K | x^2 + y^2 = 1 + dx^2y^2\}$$

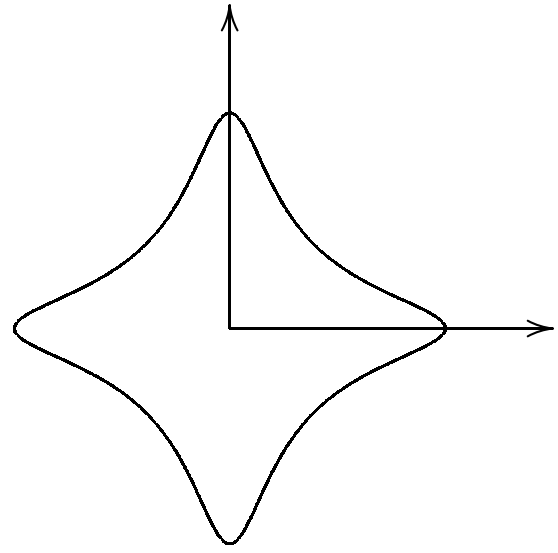

Harold M. Edwards,
(Bulletin of the AMS, **44**, 393–422, 2007)

Associative operation on points
$(x_1, y_1) + (x_2, y_2) = (x_3, y_3)$
defined by Edwards addition law

$$x_3 = \frac{x_1y_2 + y_1x_2}{1 + dx_1x_2y_1y_2} \text{ and } y_3 = \frac{y_1y_2 - x_1x_2}{1 - dx_1x_2y_1y_2}.$$

- Neutral element is $(0, 1)$.
- $-(x_1, y_1) = (-x_1, y_1)$.

Explicit formulas: addition

- $(x_1, y_1) + (x_2, y_2) = \left(\dfrac{x_1y_2 + y_1x_2}{1 + dx_1x_2y_1y_2}, \dfrac{y_1y_2 - x_1x_2}{1 - dx_1x_2y_1y_2} \right).$
- Avoid inversions: Use $(X_1 : Y_1 : Z_1)$ with $Z_1 \neq 0$ to represent $(x_1, y_1) = (X_1/Z_1, Y_1/Z_1)$, i. e., $(X_1 : Y_1 : Z_1) = (\lambda X_1 : \lambda Y_1 : \lambda Z_1)$ for $\lambda \neq 0$.
- Addition formulas in projective coordinates:

$$\begin{aligned} A &= Z_1 \cdot Z_2;\ B = A^2;\ C = X_1 \cdot X_2;\ D = Y_1 \cdot Y_2; \\ E &= d \cdot C \cdot D;\ F = B - E;\ G = B + E; \\ X_3 &= A \cdot F \cdot ((X_1 + Y_1) \cdot (X_2 + Y_2) - C - D); \\ Y_3 &= A \cdot G \cdot (D - C);\ Z_3 = F \cdot G. \end{aligned}$$

- Needs 10M + 1S + 1D + 7A, i.e. 10 multiplications, 1 squaring, 1 multiplication involving d, and 7 additions.
- Extended coordinates $(X : Y : Z : XY/Z)$

More benefits

- Doubling formulas in projective coordinates:

$$\begin{aligned} B &= (X_1+Y_1)^2;\ C=X_1^2;\ D=Y_1^2; \\ E &= C+D;\ H=Z_1^2;\ J=E-2H; \\ X_3 &= (B-E)\cdot J;\ Y_3=E\cdot(C-D);\ Z_3=E\cdot J. \end{aligned}$$

- Needs 3M + 4S + 6A.
- Does not benefit from extended coordinates, so output XY/Z only if next operation is addition; skip this coordinate as output of addition if followed by doubling.
- Precomputed points get stored as $(Y-X, Y+X, 2dXY)$; this corresponds to affine points $Z=1$ and saves the operations that depend only on one input point.

Beyond speed

- Neutral element of addition law is affine point.
- Addition law is symmetric in both inputs.
- $P+Q=\left(\frac{x_1y_2+y_1x_2}{1+dx_1x_2y_1y_2},\frac{y_1y_2-x_1x_2}{1-dx_1x_2y_1y_2}\right).$
- $[2]P=\left(\frac{x_1y_1+y_1x_1}{1+dx_1x_1y_1y_1},\frac{y_1y_1-x_1x_1}{1-dx_1x_1y_1y_1}\right).$
  Can compute this using the same formulas as general addition.
- If d is not a square in K the denominators $1+dx_1x_2y_1y_2$ and $1-dx_1x_2y_1y_2$ are never 0; addition law is complete.
- Having addition law work for doubling removes some checks from the code; addition law also works for adding $P+(-P)$ or the neutral element.