Solving covering problems with a heuristic for the p-median problem

AIRO 2005: Conference of the Italian OR Society
Camerino, Italy
September 6 to 9, 2005

Mauricio G. C. Resende

AT&T Labs Research
Florham Park, New Jersey
mgcr@research.att.com
www.research.att.com/~mgcr

Joint work with D. V. Andrade and R. F. Werneck

Summary of talk

- Heuristic for p-median problem
- Solving uncapacitated facility location problems
- Maximum covering as a p-median problem
- Set covering as a facility location problem
- Computational results

Summary of talk

- Heuristic for p-median problem
- Solving uncapacitated facility location problems
- Maximum covering as a p-median problem
- Set covering as a facility location problem
- Computational results

AT&T

Summary of talk

- Heuristic for p-median problem
- Solving uncapacitated facility location problems
- Maximum covering as a p-median problem
- Set covering as a facility location problem
- Computational results

AT&T

Summary of talk

- Heuristic for p-median problem
- Solving uncapacitated facility location problems
- Maximum covering as a p-median problem
- Set covering as a facility location problem
- Computational results

Summary of talk

- Heuristic for p-median problem
- Solving uncapacitated facility location problems
- Maximum covering as a p-median problem
- Set covering as a facility location problem
- Computational results

AT&T

Facility location

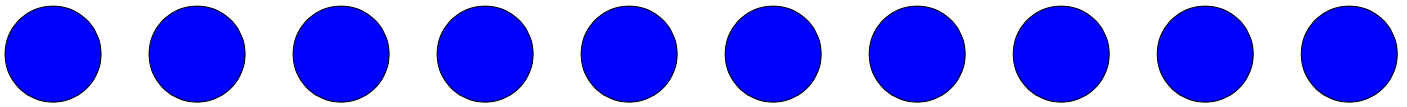

Facility location

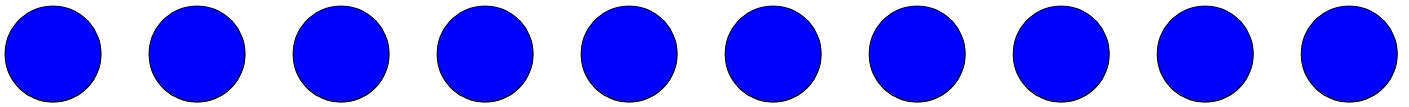

facilities

Facility location

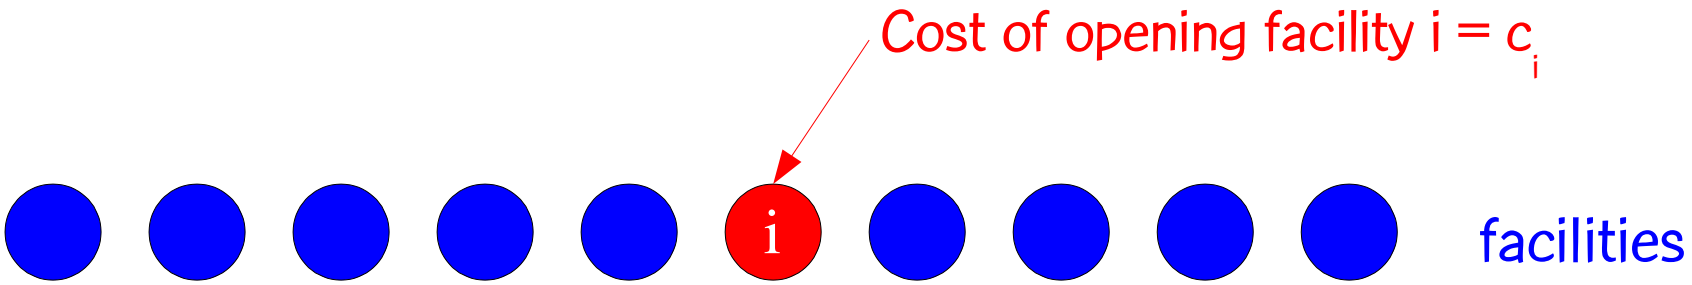

Facility location

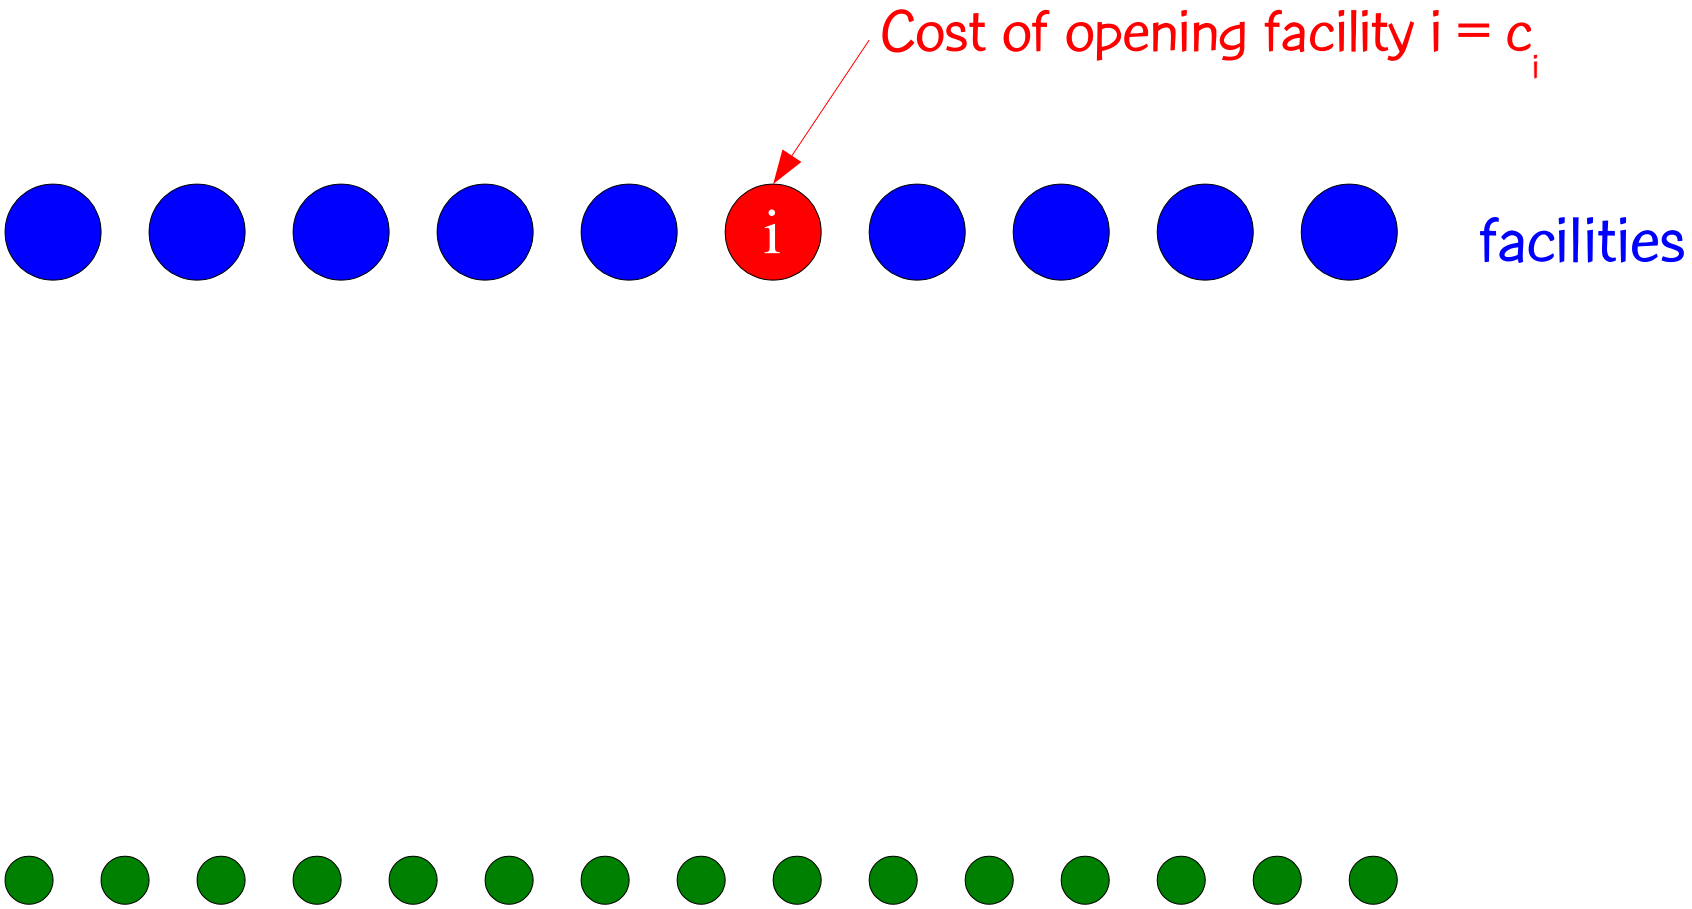

Facility location

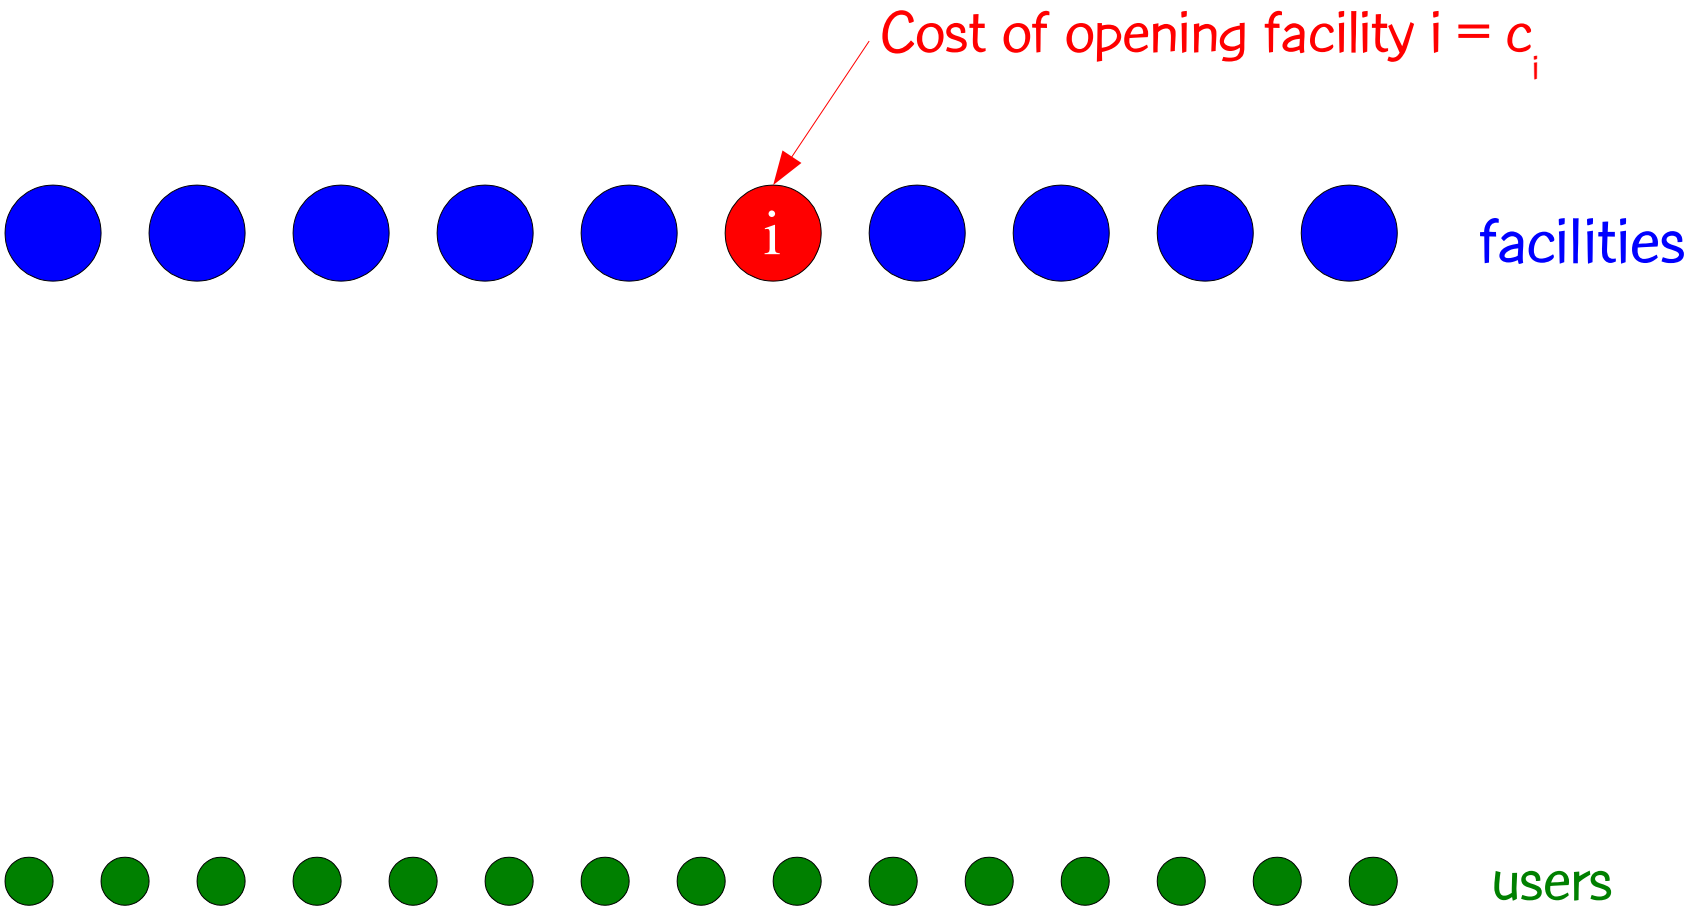

Facility location

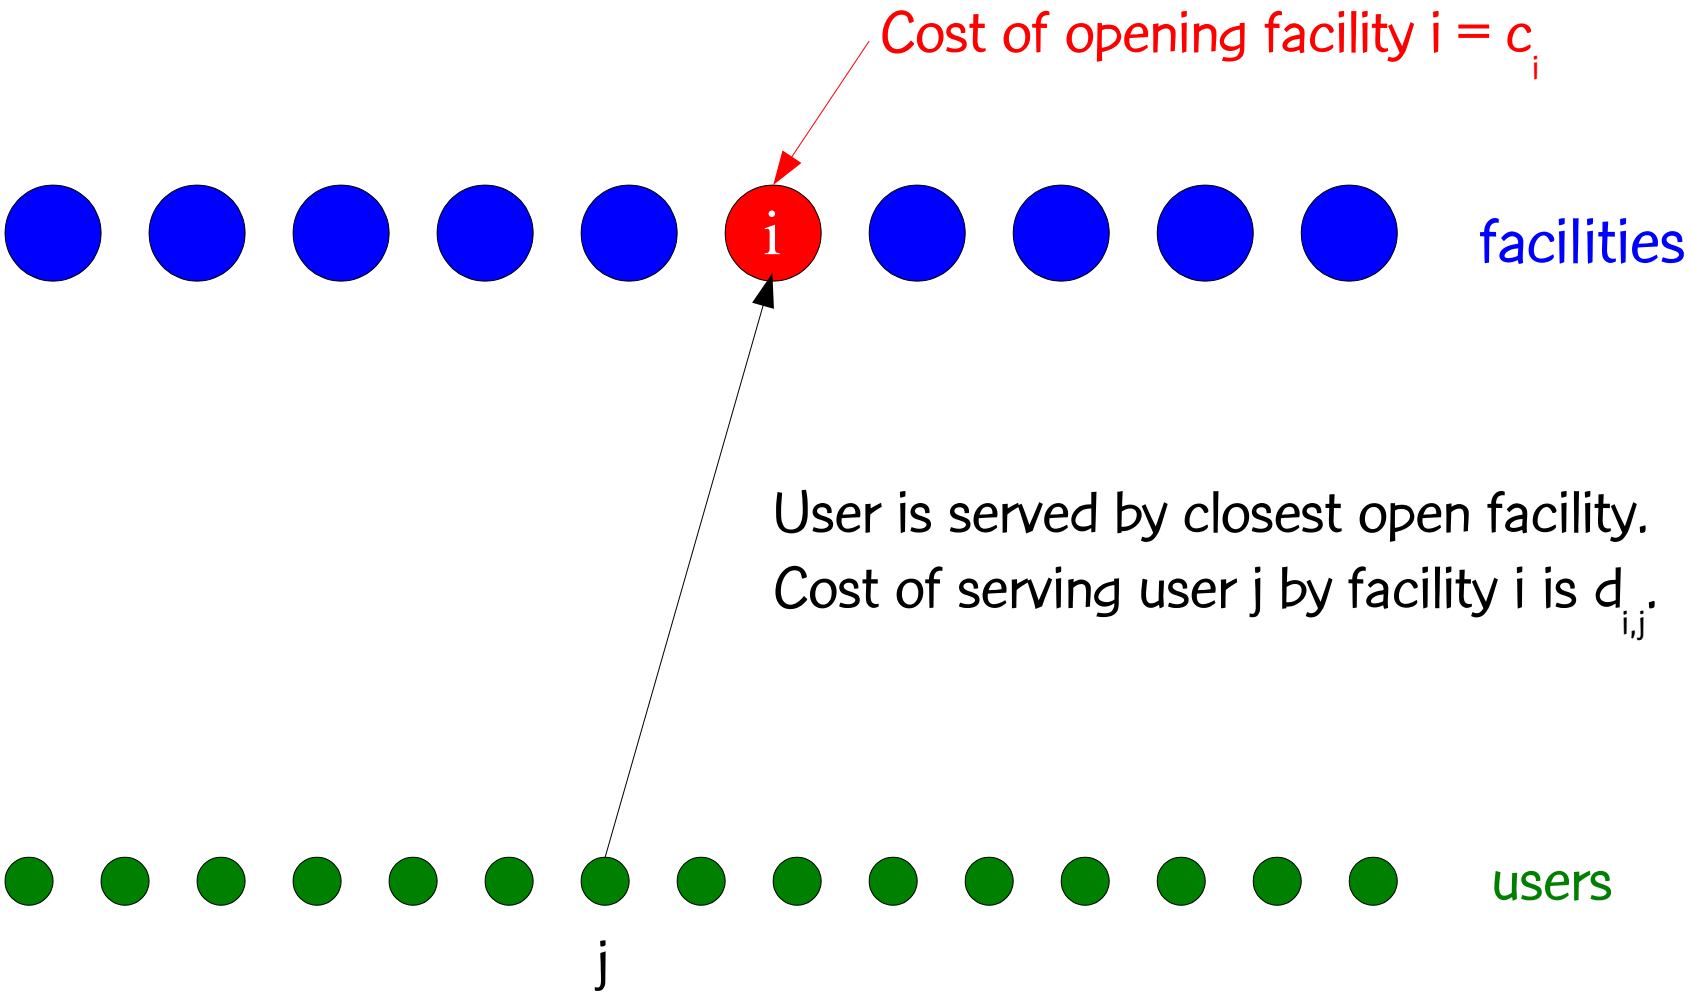

p-median problem

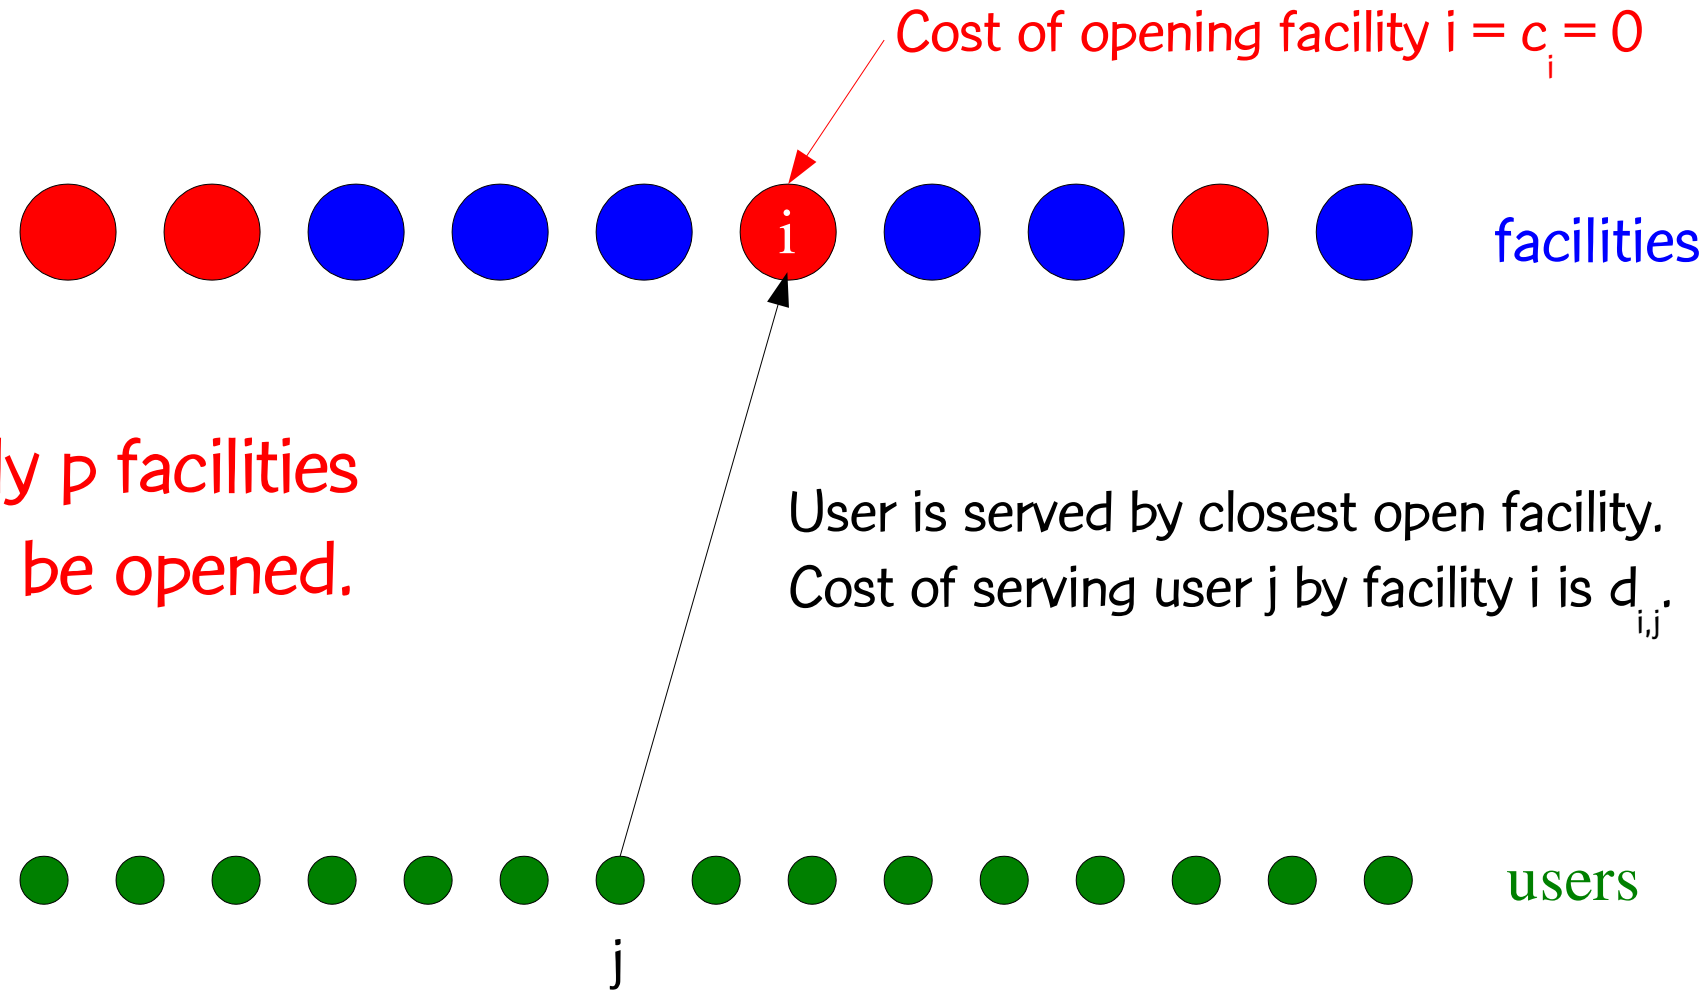

AT&T

GRASP with path-relinking for p-median

- Resende & Werneck (2003) describe a multi-start heuristic for p-median
- Each iteration:
  - Construct greedy randomized solution
  - Apply fast version of Whitaker's swap-based local search
  - Apply path-relinking between local optimum and a randomly chosen solution from set of elite solutions
- Every k iterations, intensify elite set by performing path-relinking between all pairs of elite solutions

GRASP with path-relinking for facility location

- Resende & Werneck (2004) extended multi-start heuristic for p-median to solve uncapacitated facility location problems
- At iteration i, we determine the number p_i of m facilities to open and apply p-median heuristic for $p = p_i$
  - For $i = 1$, $p_i = \lceil m/2 \rceil$;
  - For $i > 1$, we pick the average number of facilities opened in the first $i - 1$ iterations;
- In local search and PR, allow opening or closing facilities in addition to usual swapping

References

- M.G.C. Resende and R.F. Werneck, "A fast swap-based local search procedure for location problems," AT&T Labs Research Technical Report TD-5R3KBH, Florham Park, NJ, Sept. 2003, revised Dec. 2004. To appear in Annals of Operations Research.

`http://www.research.att.com/~mgcr/doc/locationls.pdf`

References

- M.G.C. Resende and R.F. Werneck, "A hybrid heuristic for the p-median problem," Journal of Heuristics, vol. 10, pp. 59–88, 2004.

`http://www.research.att.com/~mgcr/doc/hhpmedian.pdf`

References

- M.G.C. Resende and R.F. Werneck, "A hybrid multi-start heuristic for the uncapacitated facility location problem," AT&T Labs Research Technical Report TD-5RELRR, Florham Park, NJ, Sept. 2003, revised Nov. 2004. To appear in European J. of Operational Research, 2005.

  `http://www.research.att.com/~mgcr/doc/guflp.pdf`

Max covering problem

- Given:
  - a set of users U where user u has weight w_u
  - a set of facilities F, where facility f cover users $S_f \subseteq U$ if it is opened
  - a positive number p of facilties to be opened
- Objective:
  - maximize the sum of the weights of covered users

Max covering as a p-median problem

facilities

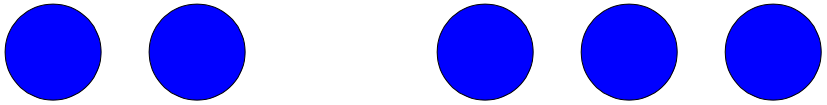

users

Max covering as a p-median problem

facilities

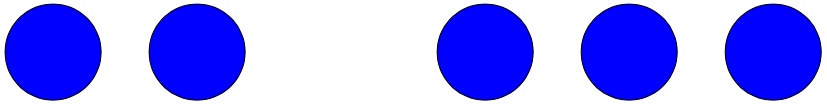

user u

AT&T

Max covering as a p-median problem

Red facilities cover user u

Black facilities do not cover user u

facilities

user u

AT&T

Max covering as a p-median problem

facilities

Red facilities cover user u

Black facilities do not cover user u

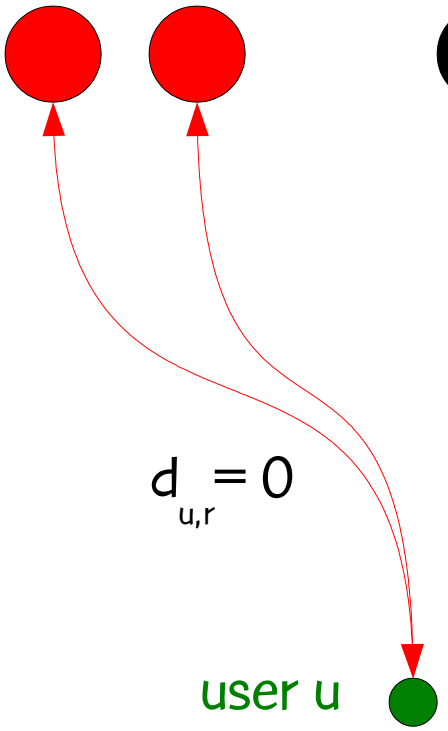

Max covering as a p-median problem

facilities

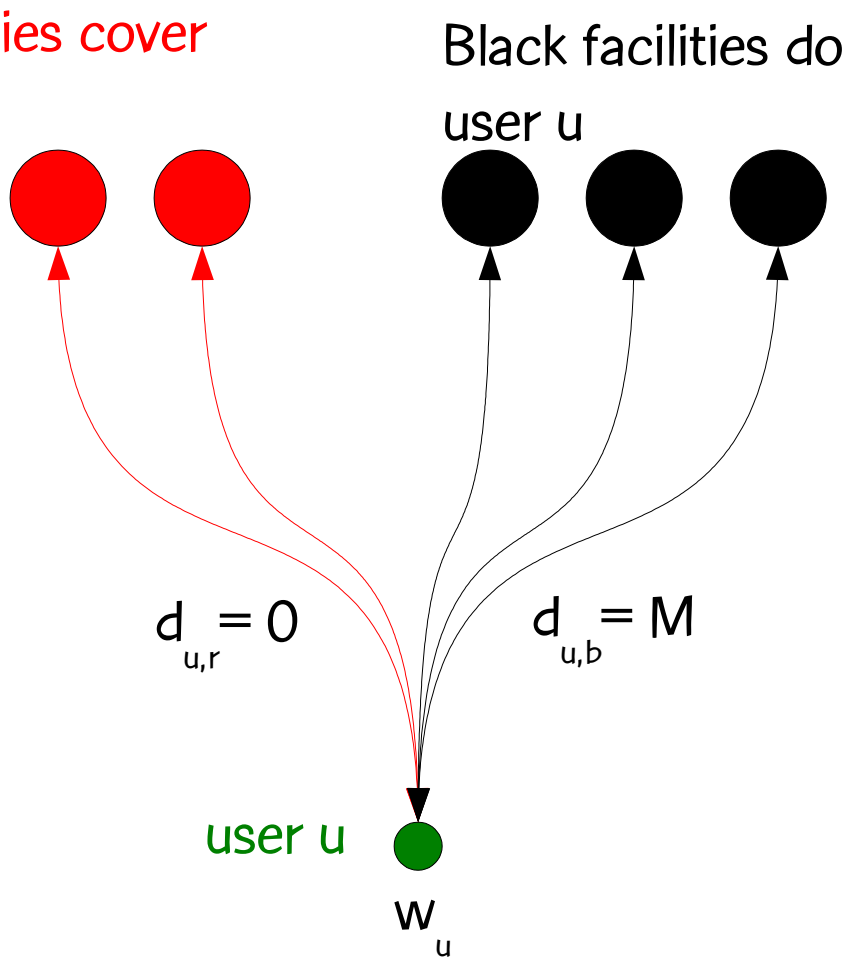

$$M = 1 + \max \{w_i \mid i \in U\}$$

AT&T

Max covering as a p-median problem

facilities

Red facilities cover user u

Black facilities do not cover user u

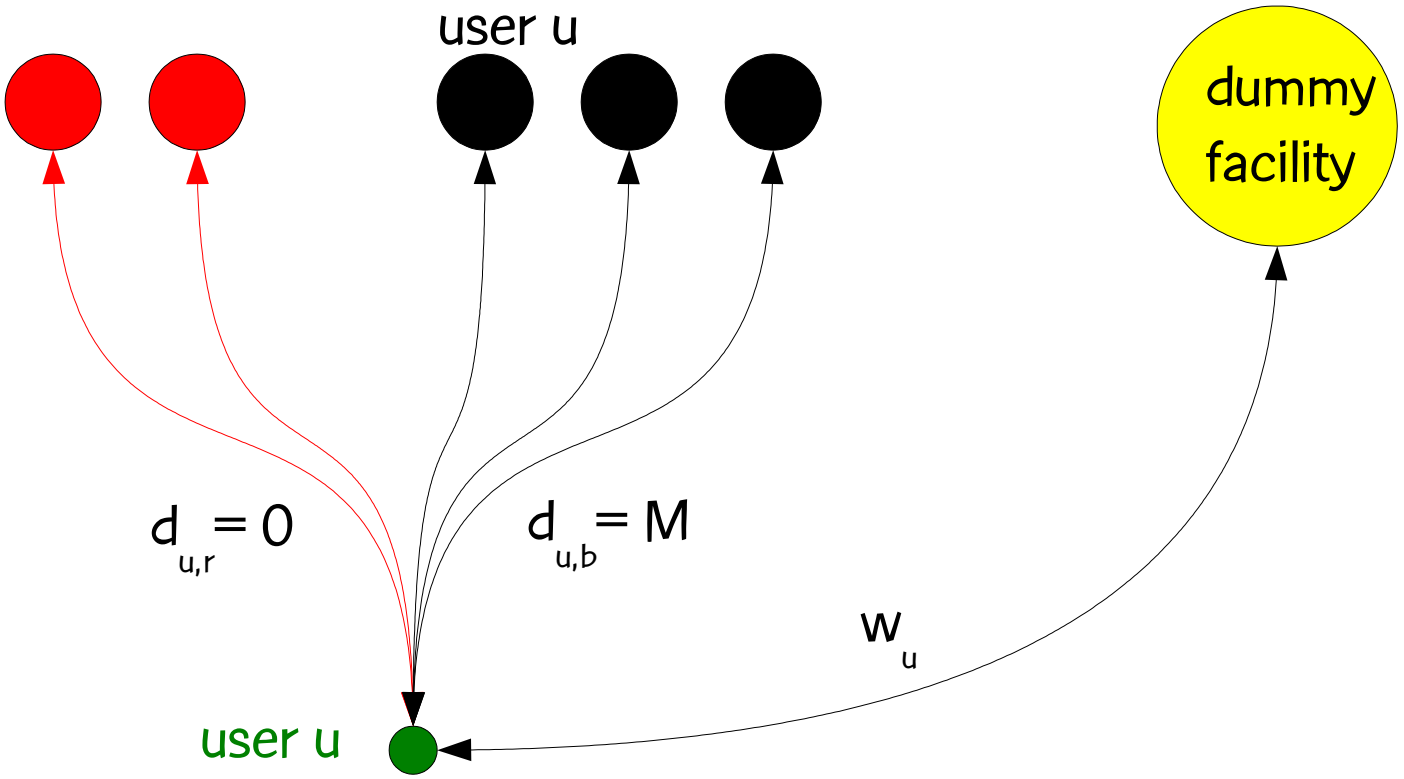

$$M = 1 + \max \{w_i \mid i \in U\}$$

AT&T

Max covering as a p-median problem

facilities

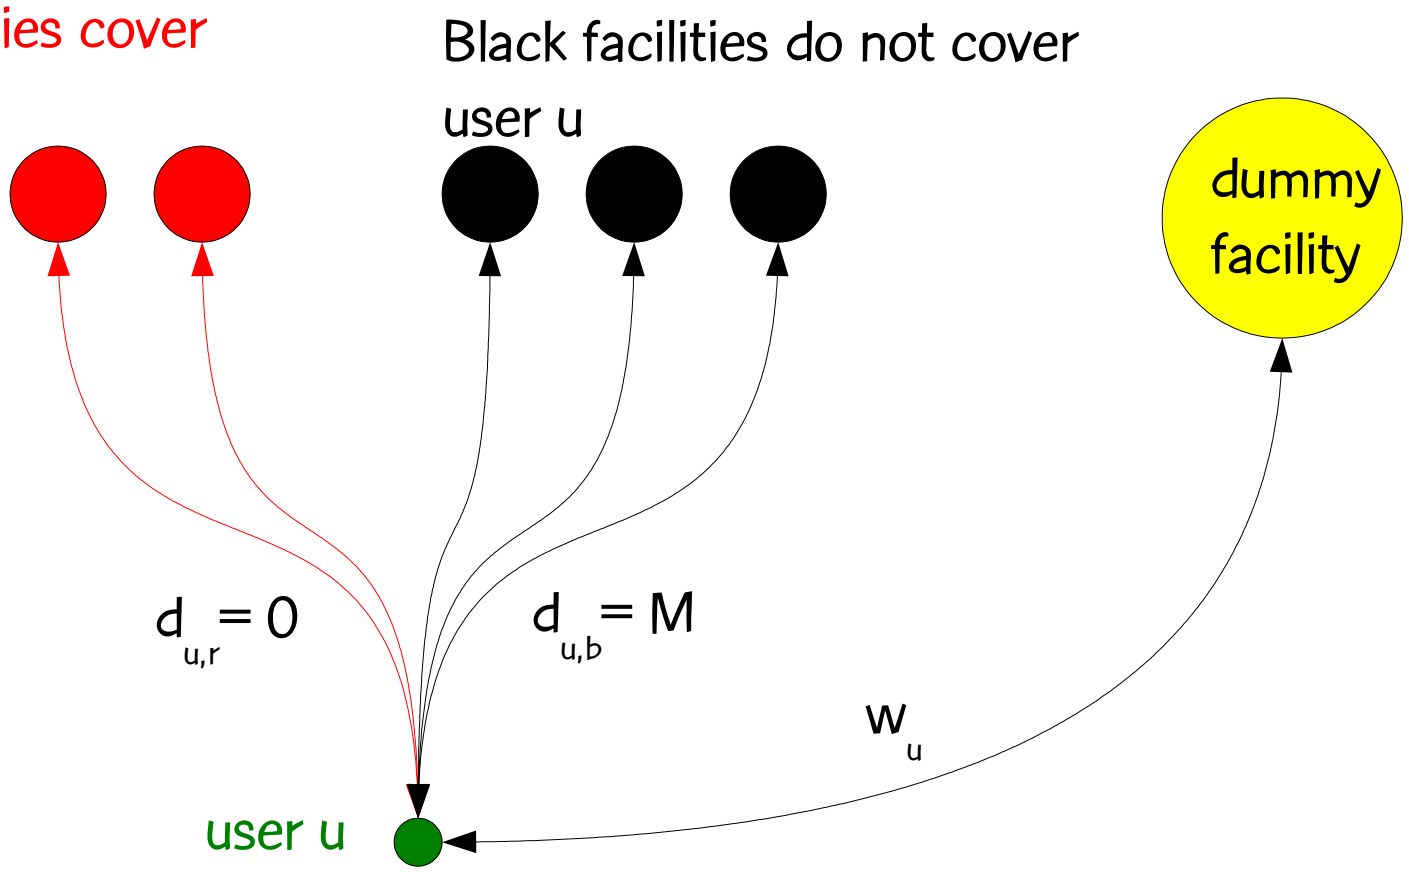

$M = 1 + \max \{w_i \mid i \in U\}$

If no red facility is open, user u is covered by dummy

AT&T

Max covering as a p-median problem

facilities

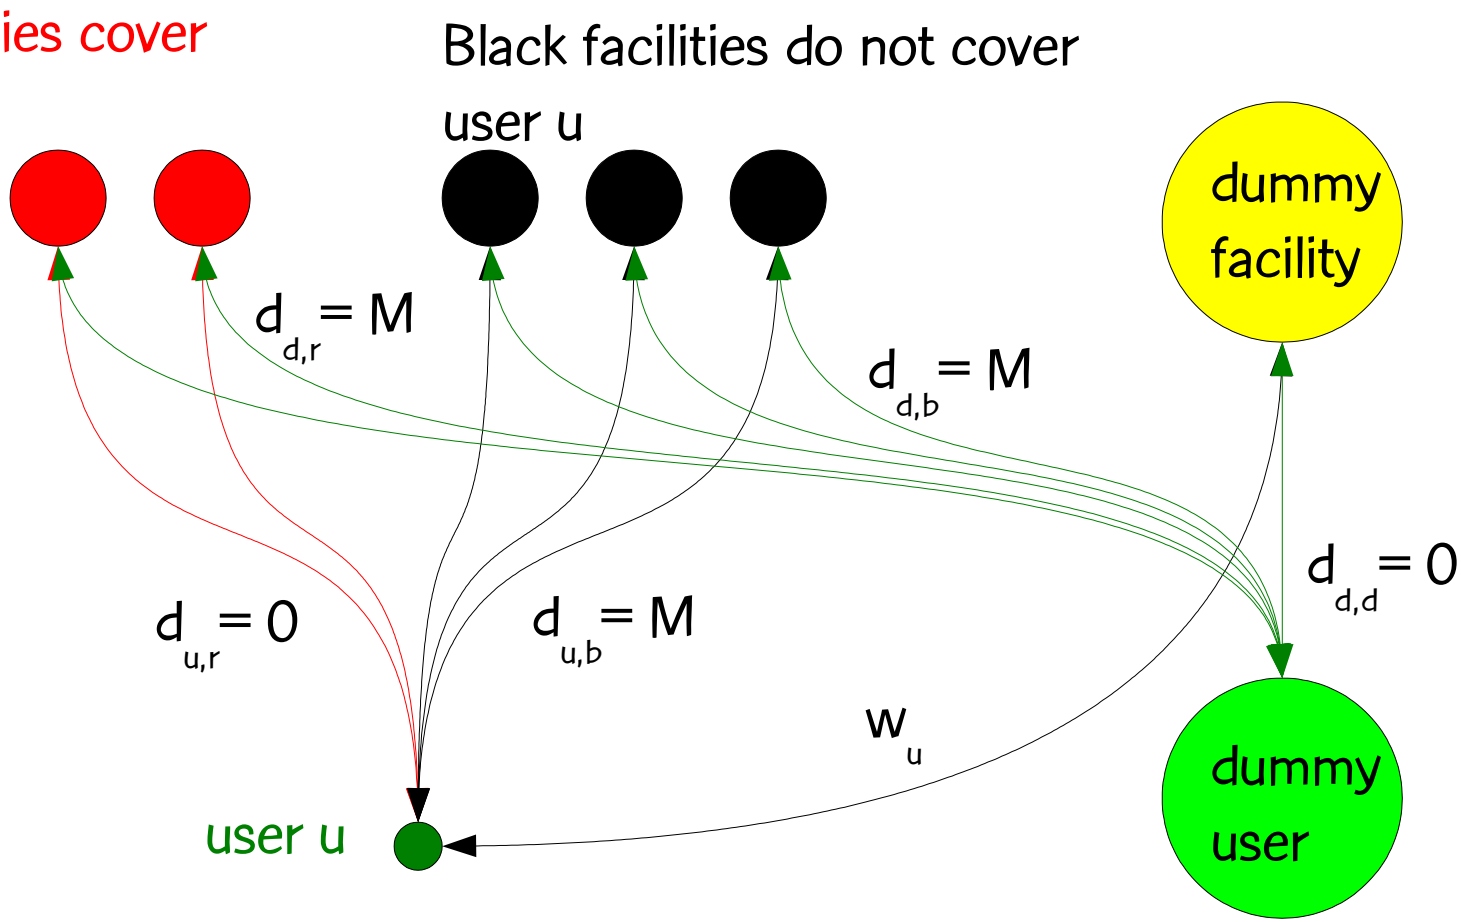

$M = 1 + \max \{w_i \mid i \in U\}$

Dummy user insures that dummy facility is open.
P+1 facilities are to be opened.

AT&T

Max covering as a p-median problem

Red facilities cover user u

Black facilities do not cover user u

facilities

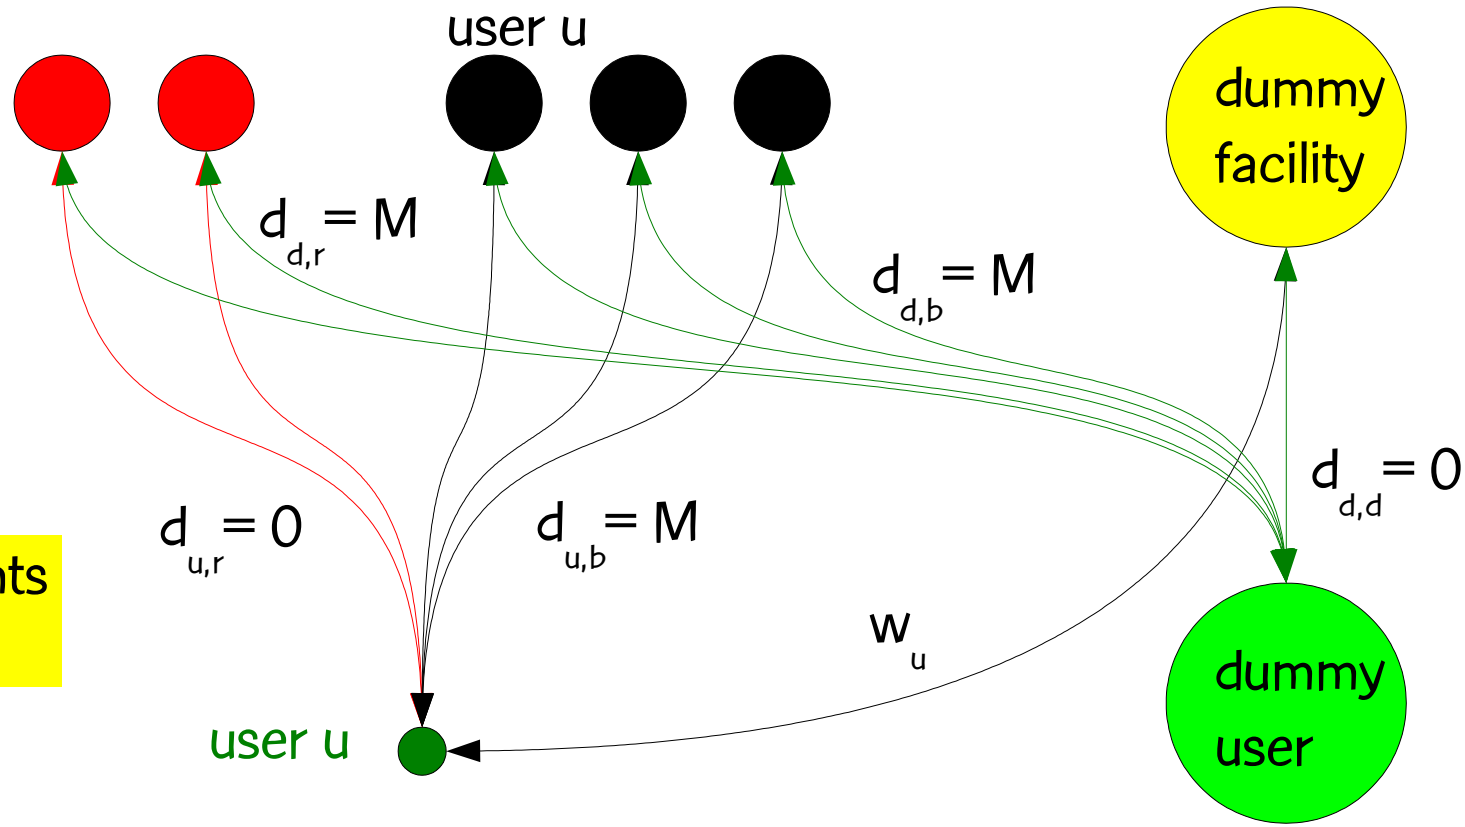

Minimize sum of weights of uncovered users.

$$M = 1 + \max \{w_i \mid i \in U\}$$

Dummy user insures that dummy facility is open.
P+1 facilities are to be opened.

AT&T

Modem pool placement for Internet service provider

- Worldnet: AT&T's Internet Service Provider
- Dial-up: hundreds of points of presence (PoPs)
  - Telephone numbers customers must call when making an Internet connection.
- Current footprint:
  - 1305 PoPs;
  - 77.66% of the telephone numbers in the U.S. can make local calls to Worldnet.

AT&T

Worldnet

- When is a call local?
  - Not simply "within same area code".
  - Telephone system divided into exchanges:
    - Area code + first three digits (973360, for example).
- Each PoP has a coordinate.
- We know which exchanges can make local calls to each coordinate (the coverage).
  - Just a big table;
  - 69,534 exchanges covered by current footprint.

Expanding the Footprint of Worldnet

- Problem:
  - Increase coverage beyond current 77.66%.
- AT&T can use UUNet PoPs:
  - 1,498 candidate PoPs.
  - 568 of those cover at least one new exchange.
- Main question:
  - If we want to open p new PoPs, which PoPs do we open?
    - Goal: maximize coverage.
- This is the maximum cover problem.

Expansion

568 facilities (each is a potential PoP)

89 537 users (weight of each user is number of phones in exchange)

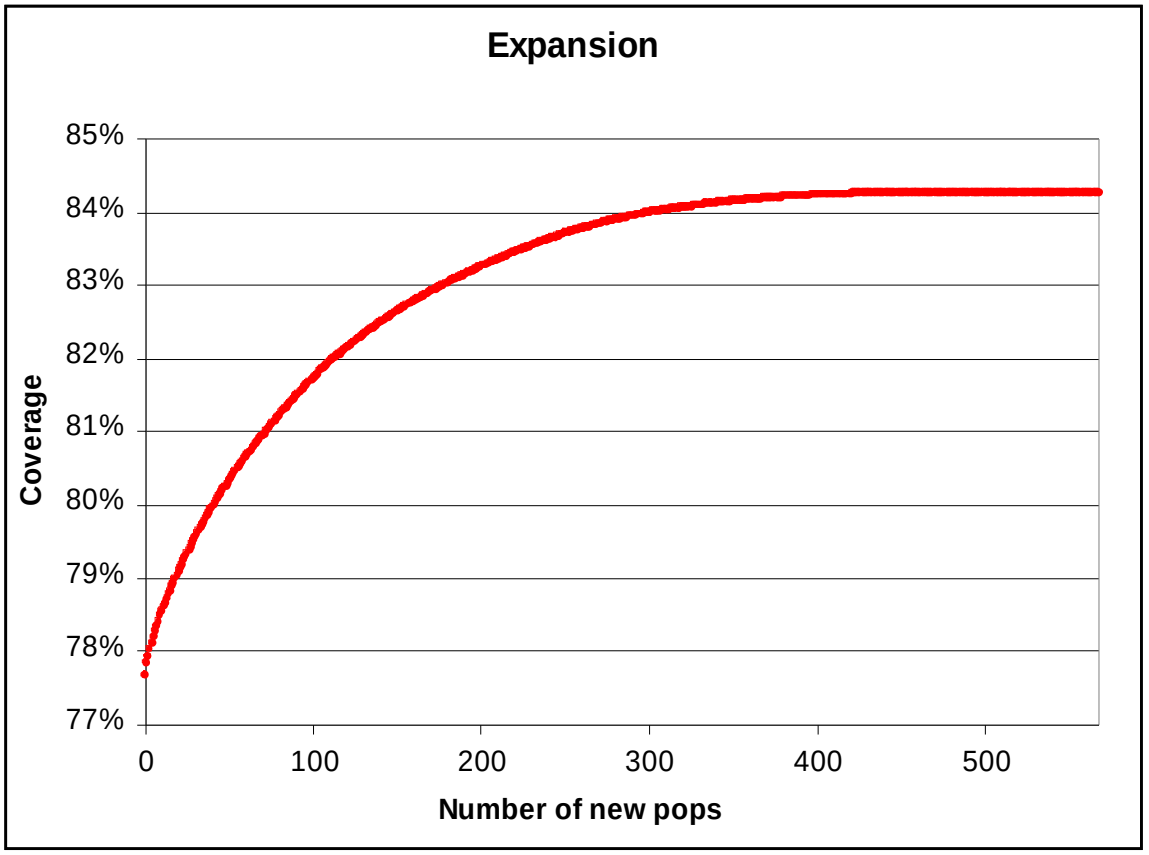

Coverage	Footprint
77.66%	current
78%	current+3
79%	current+19
80%	current+41
81%	current+72
82%	current+113
83%	current+177
84%	current+301
84.27%	current+464

AT&T

Set covering problem

- Given:
  - a set of users U
  - a set of facilities F, where facility $f \in F$ covers users $S_f \subseteq U$ if it is opened, where facility f has weight w_f
- Objective:
  - minimize the sum of the weights of opened facilities such that all users are covered by at least one opened facility

Set covering as a facility location problem

facilities

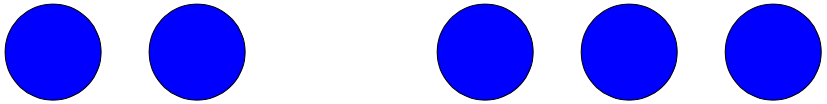

users

Set covering as a facility location problem

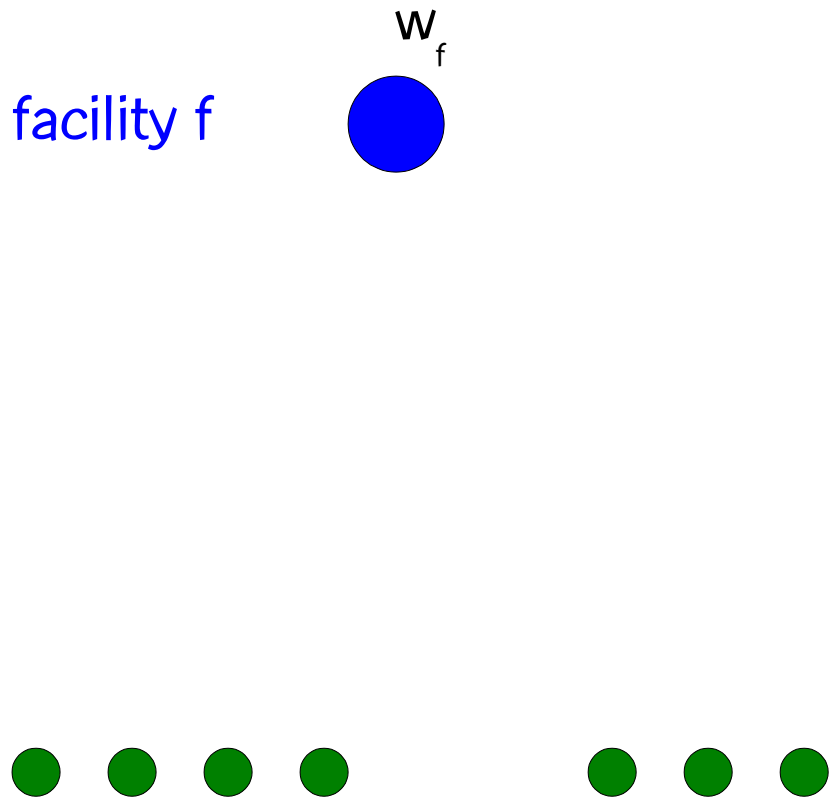

Set covering as a facility location problem

users

Red users are covered by facility f.

Black users are not covered by facility f.

Set covering as a facility location problem

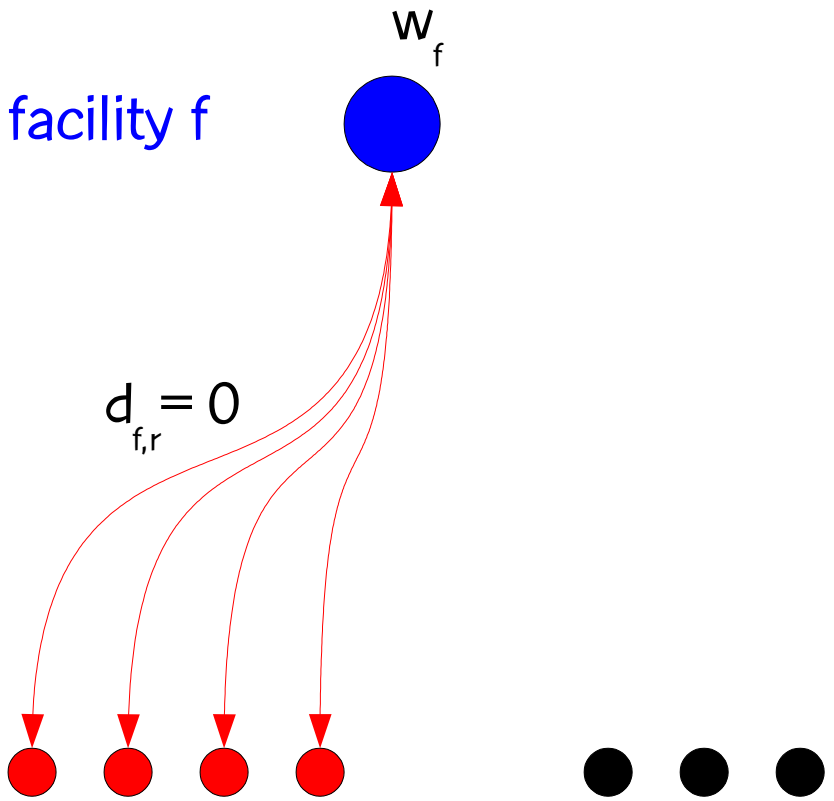

Red users are covered by facility f.

Black users are not covered by facility f.

Set covering as a facility location problem

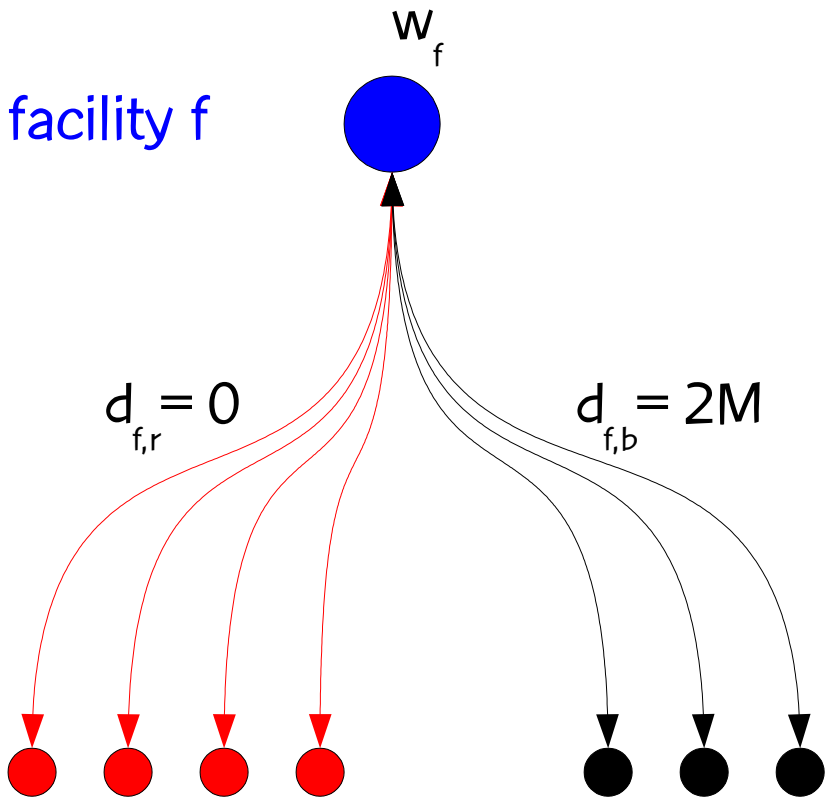

Red users are covered by facility f.

Black users are not covered by facility f.

$$M = 1 + \max \{w_f \mid f \in F\}$$

Set covering as a facility location problem

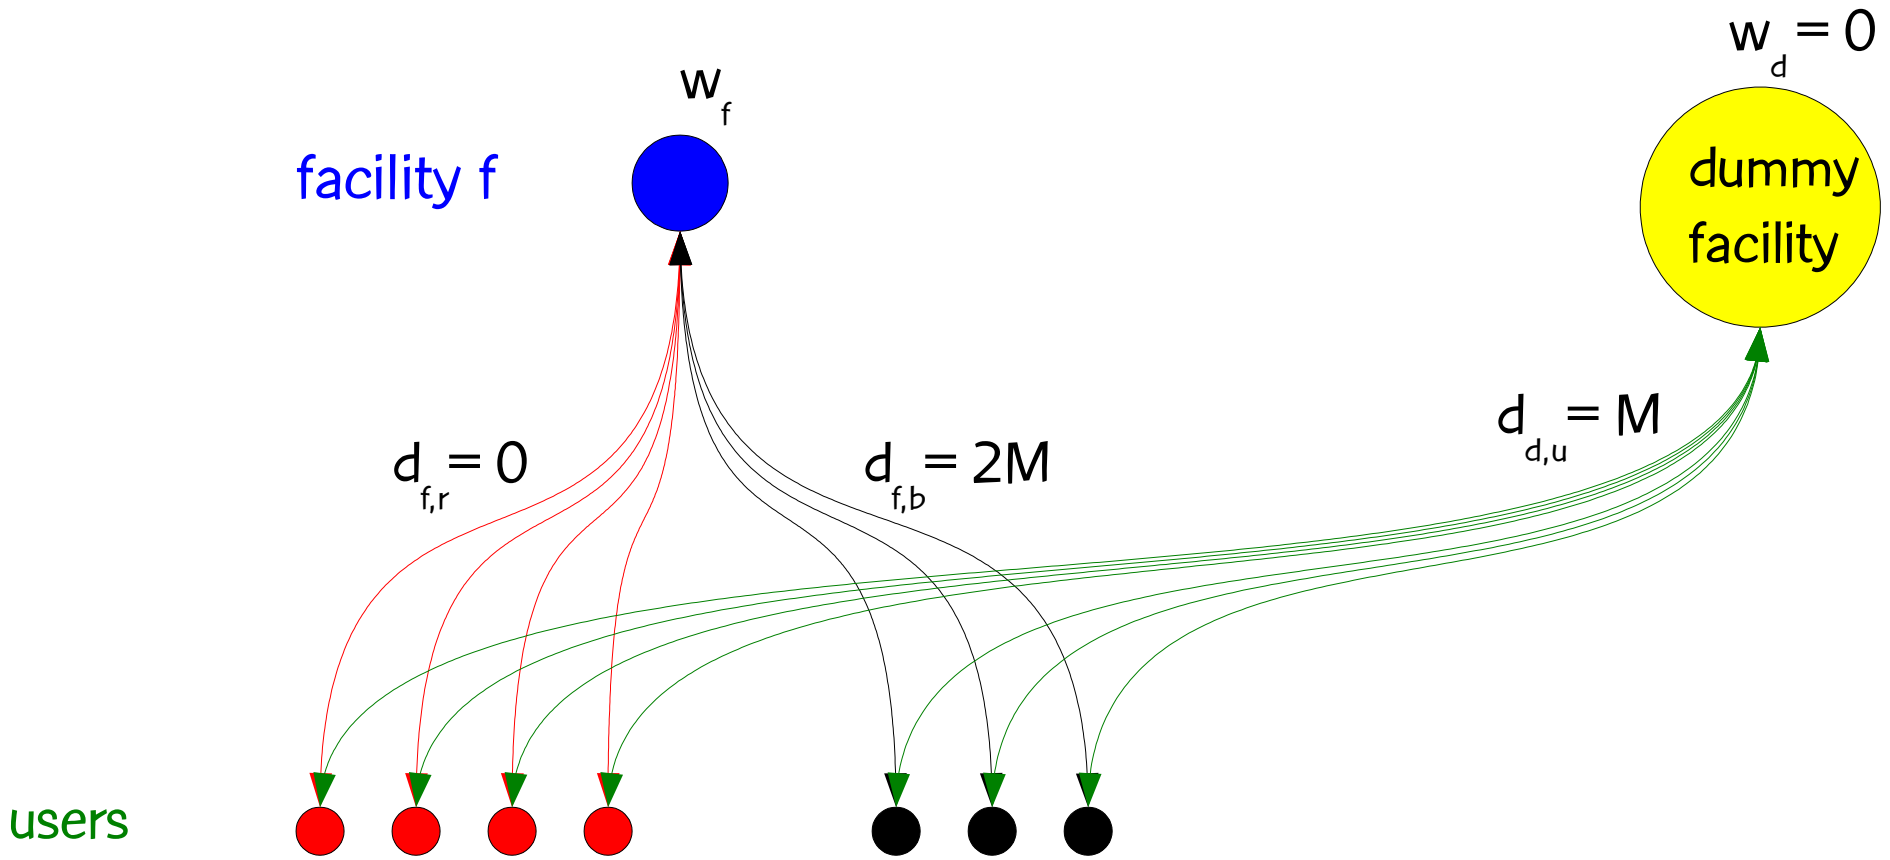

Red users are covered by facility f.

Black users are not covered by facility f.

$$M = 1 + \max \{w_f \mid f \in F\}$$

Dummy facility insures that all users are covered.

Set covering as a facility location problem

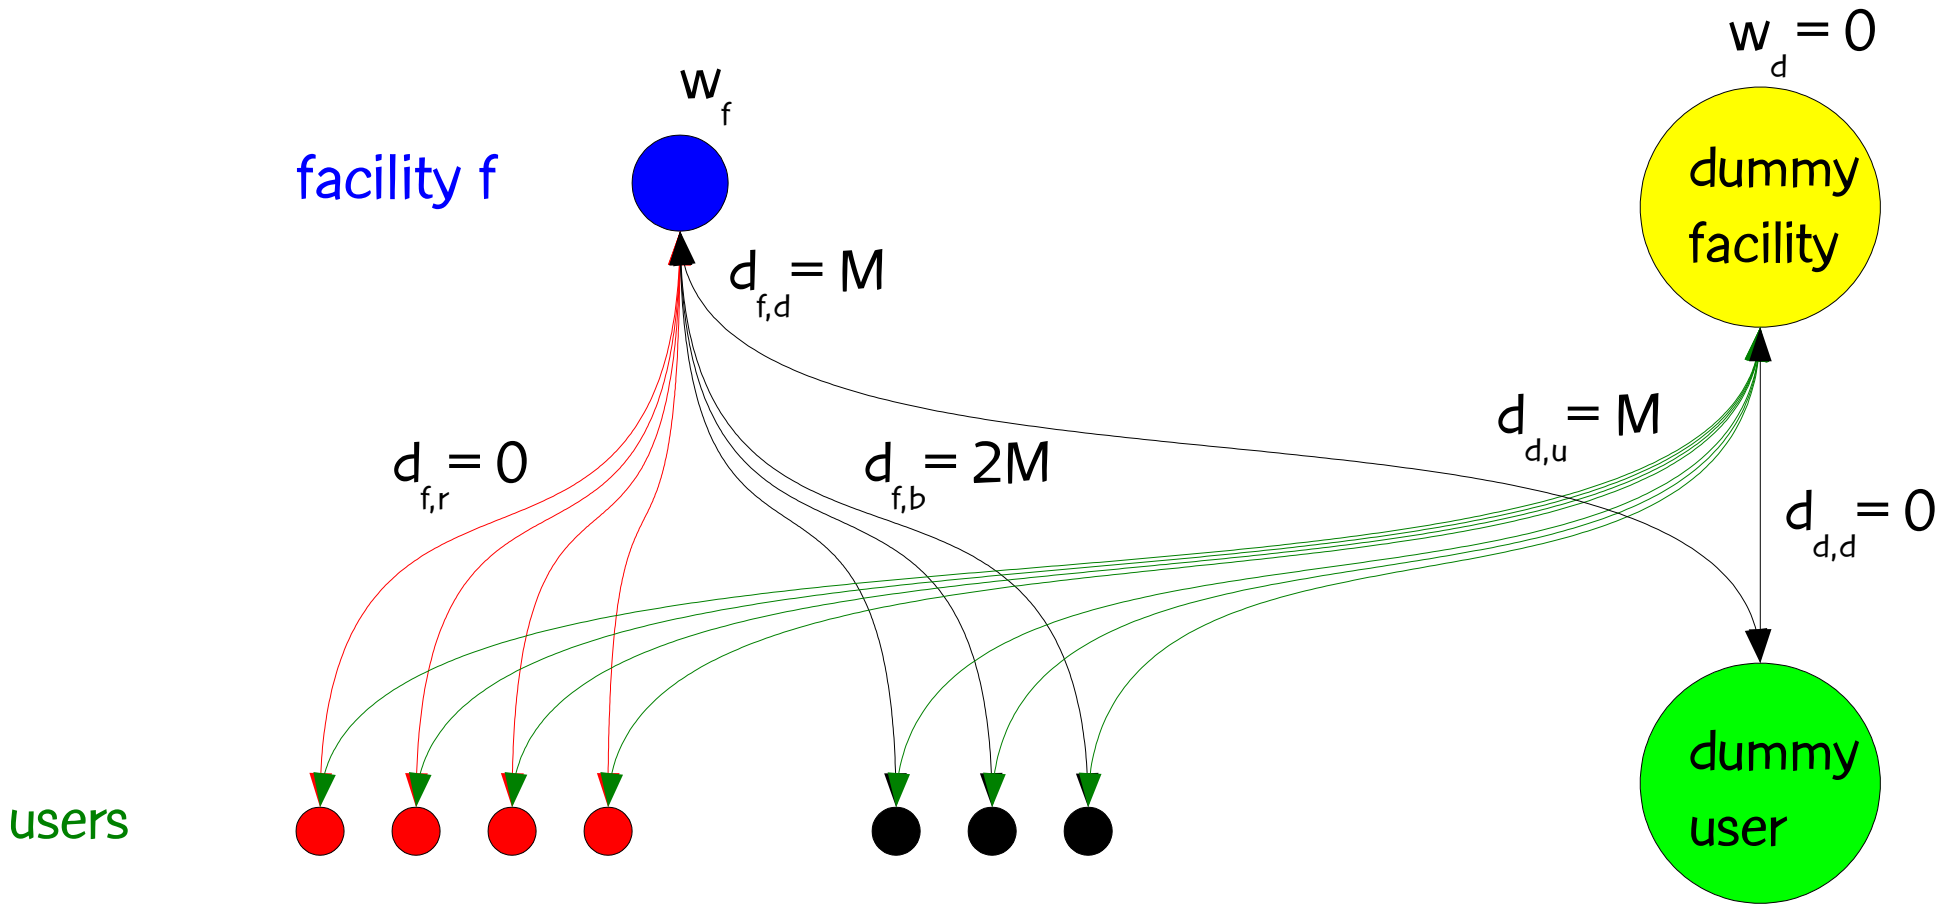

Red users are covered by facility f.

Black users are not covered by facility f.

$$M = 1 + \max \{w_f \mid f \in F\}$$

Dummy user insures that dummy facility is opened.

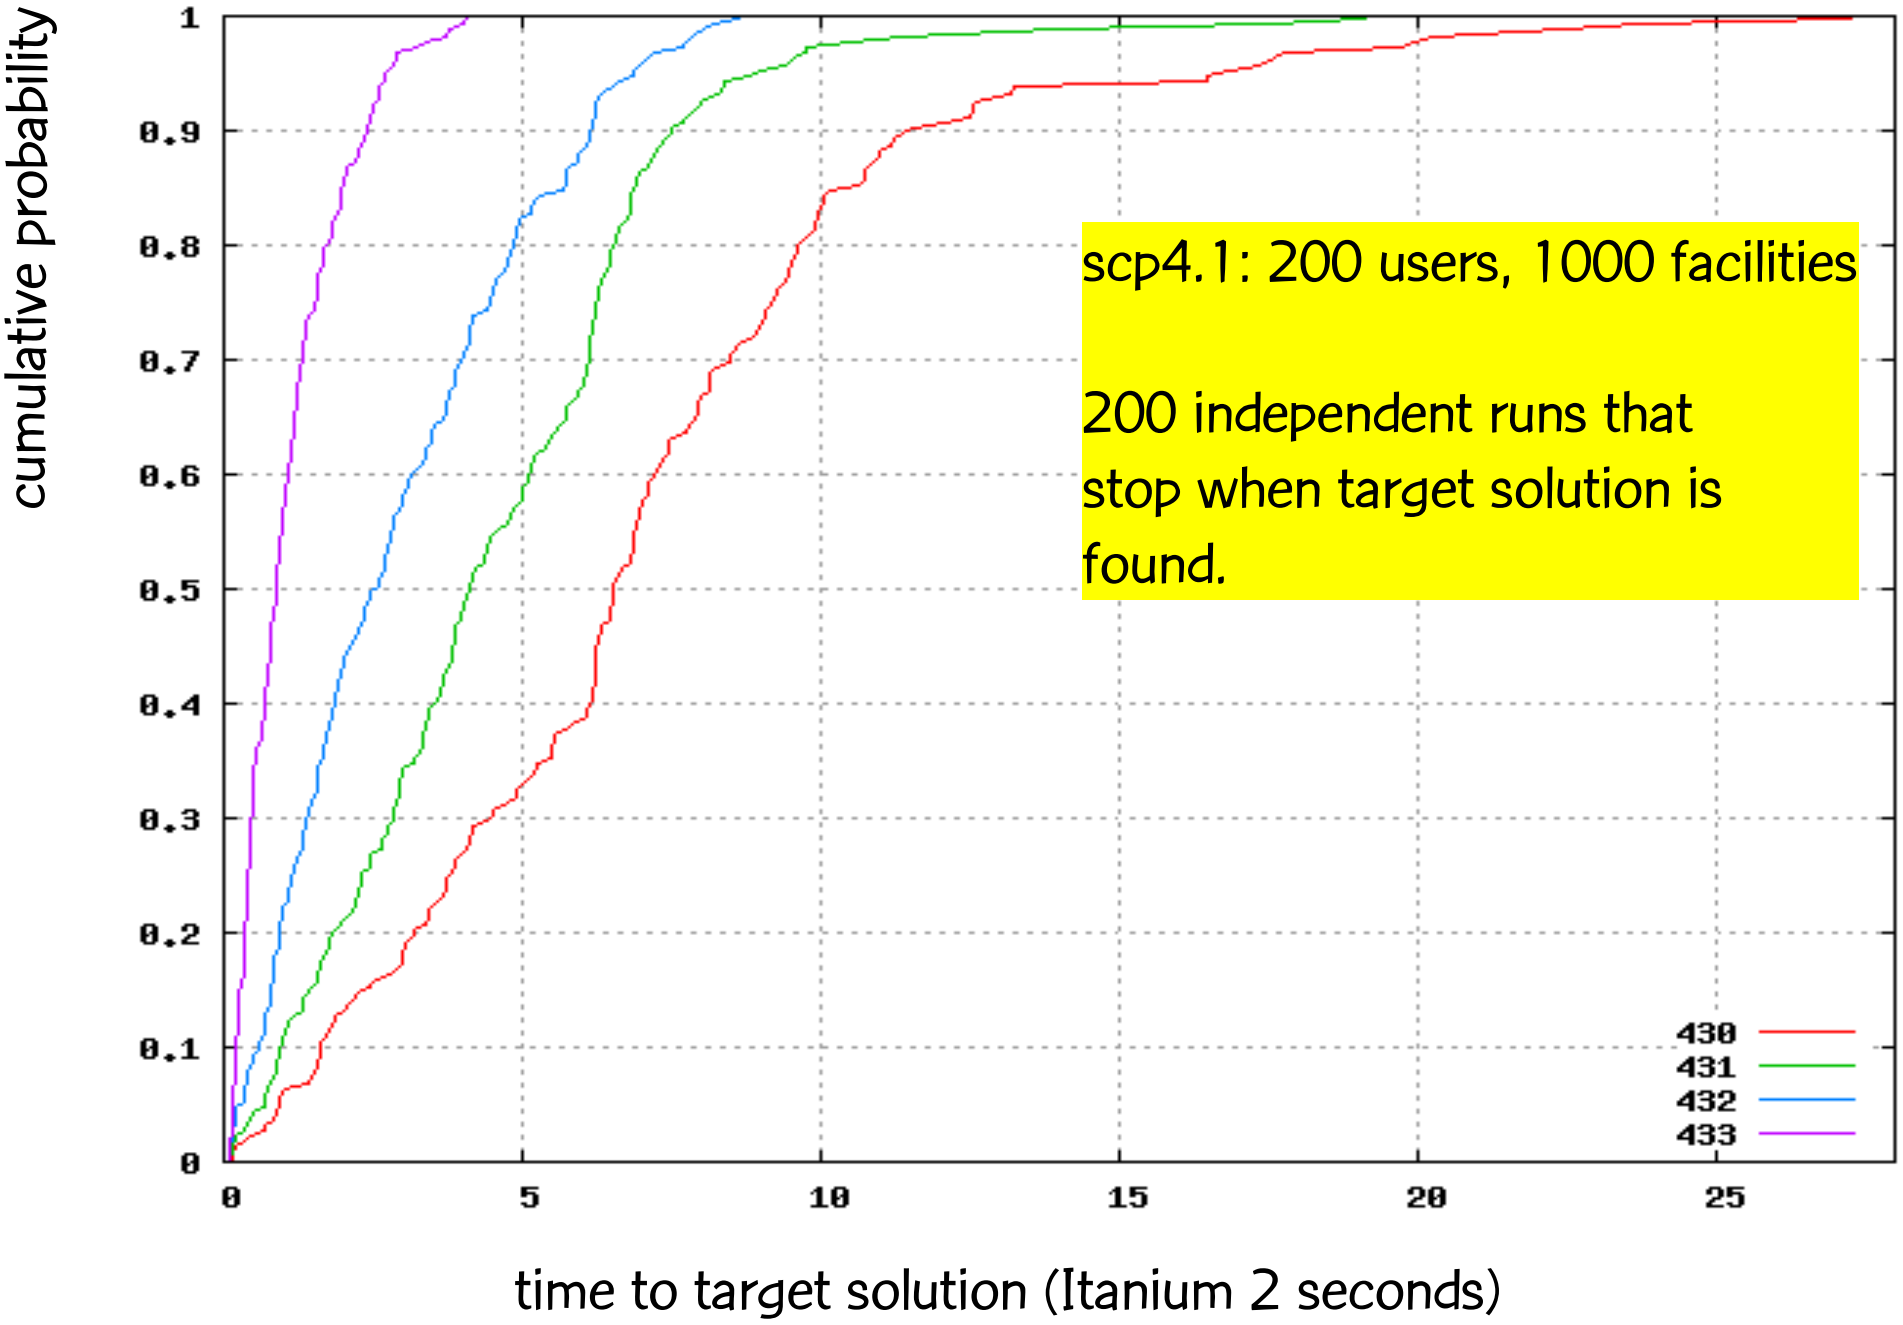
cumulative probability
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
5
10
15
20
25
time to target solution (Itanium 2 seconds)
scp4.1: 200 users, 1000 facilities
200 independent runs that stop when target solution is found.
430
431
432
433

AT&T

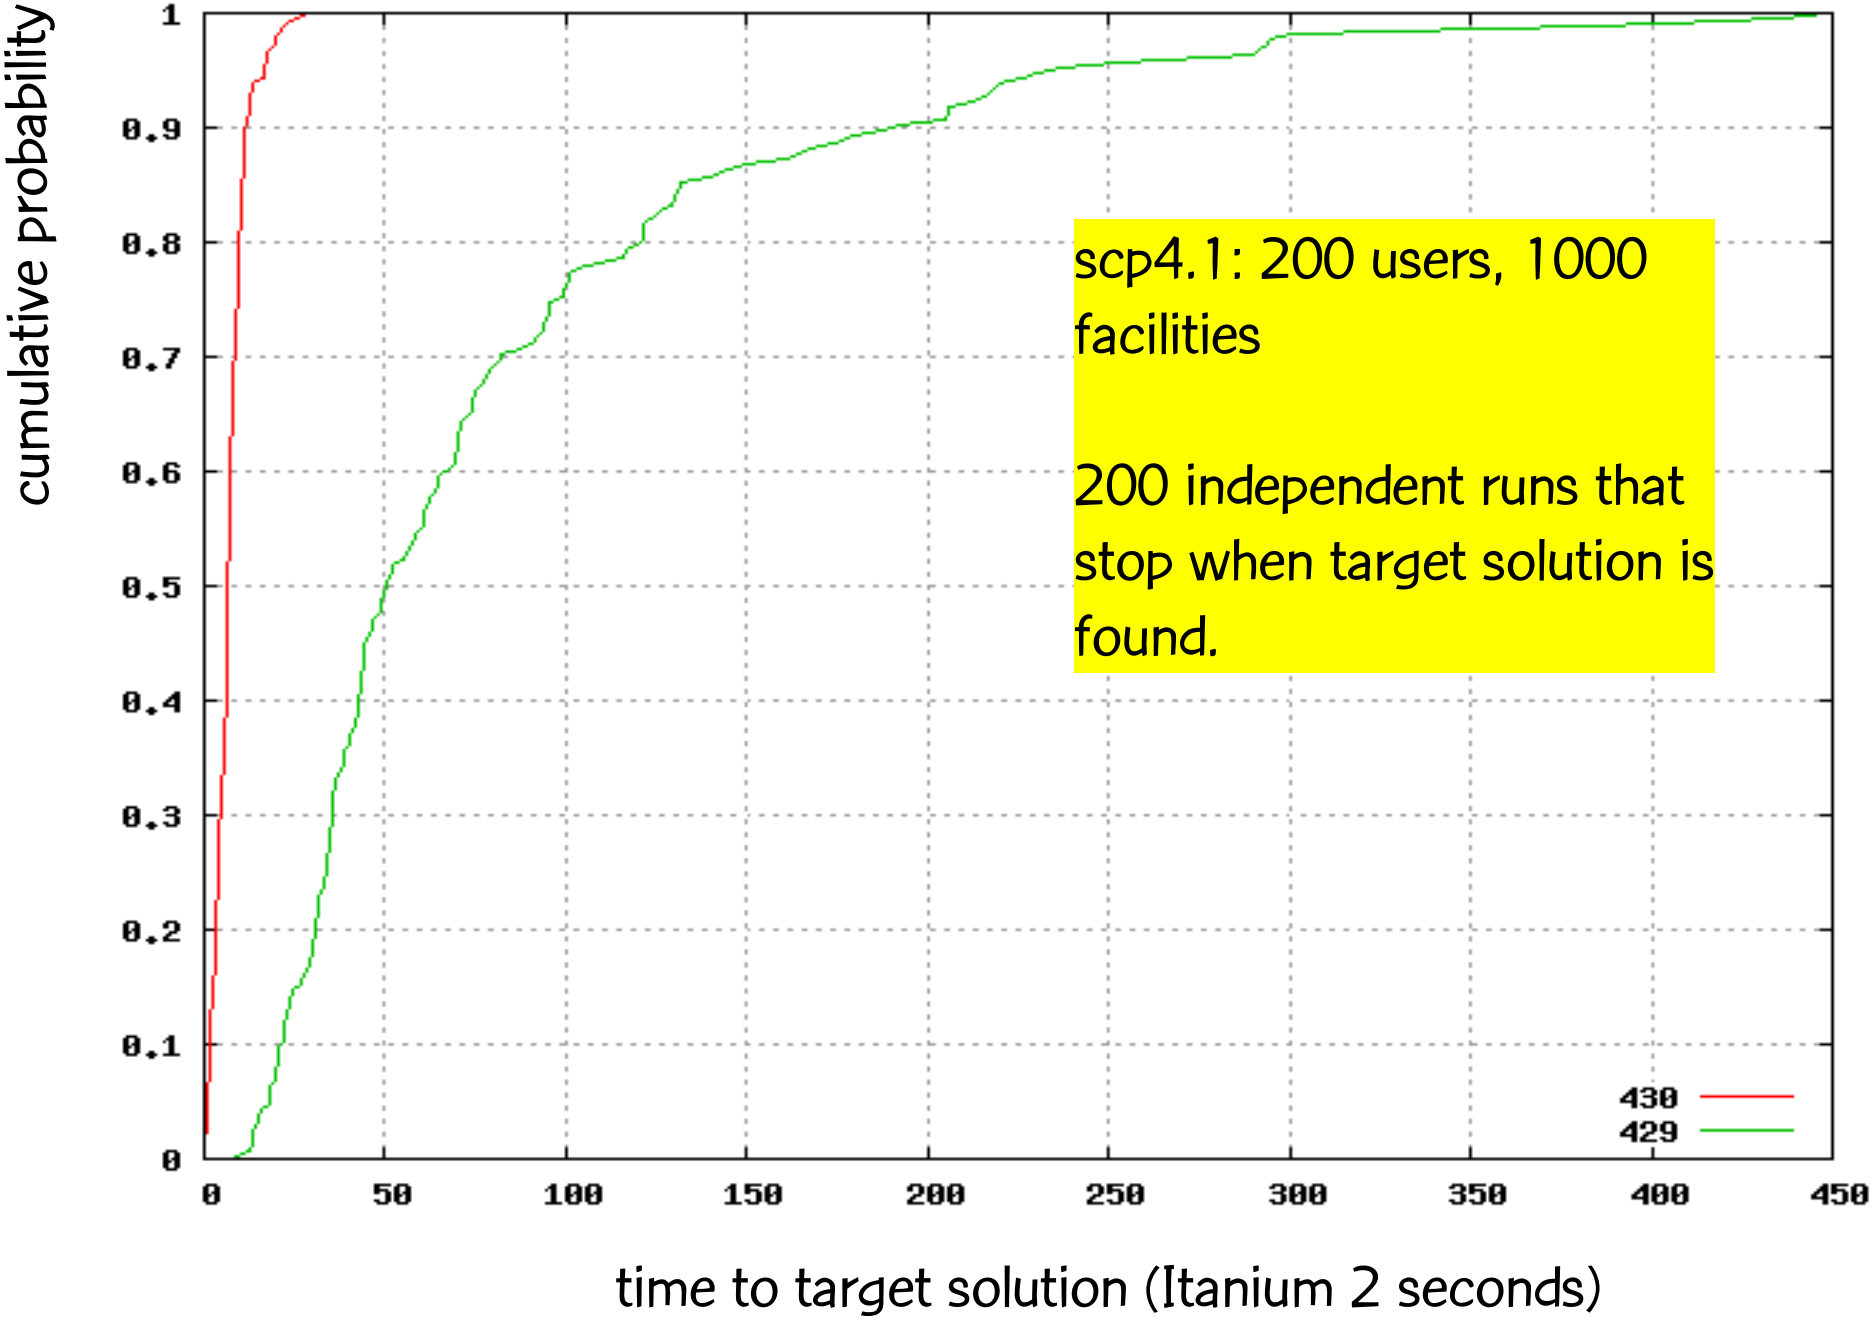
cumulative probability
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
50
100
150
200
250
300
350
400
450
time to target solution (Itanium 2 seconds)
scp4.1: 200 users, 1000 facilities
200 independent runs that stop when target solution is found.
430
429

AT&T

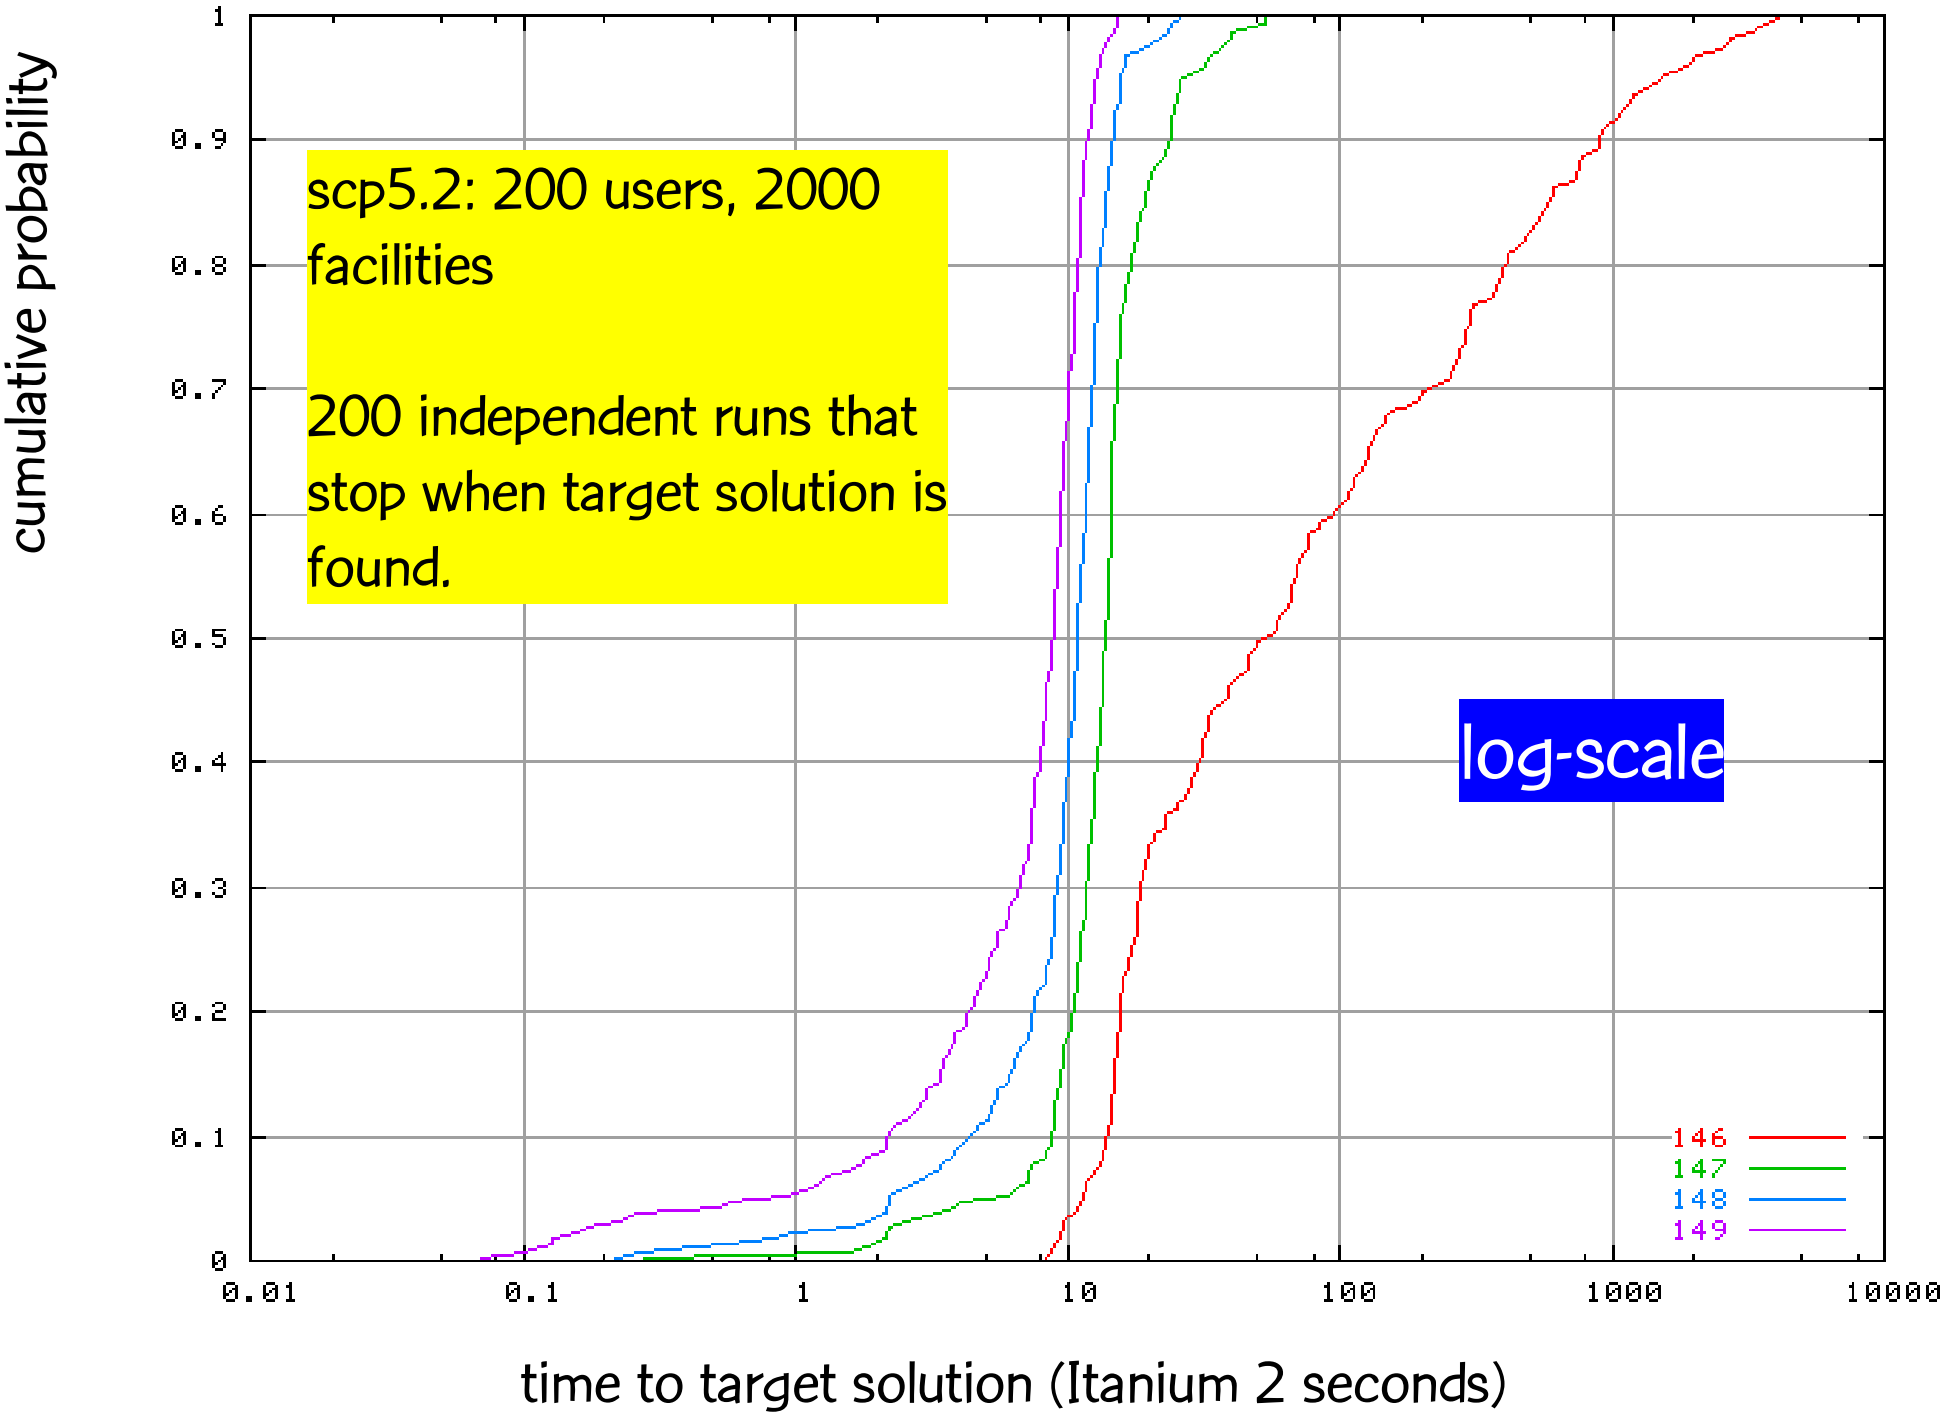
cumulative probability
scp5.2: 200 users, 2000 facilities
200 independent runs that stop when target solution is found.
log-scale
146
147
148
149
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0.01
0.1
1
10
100
1000
10000
time to target solution (Itanium 2 seconds)

AT&T

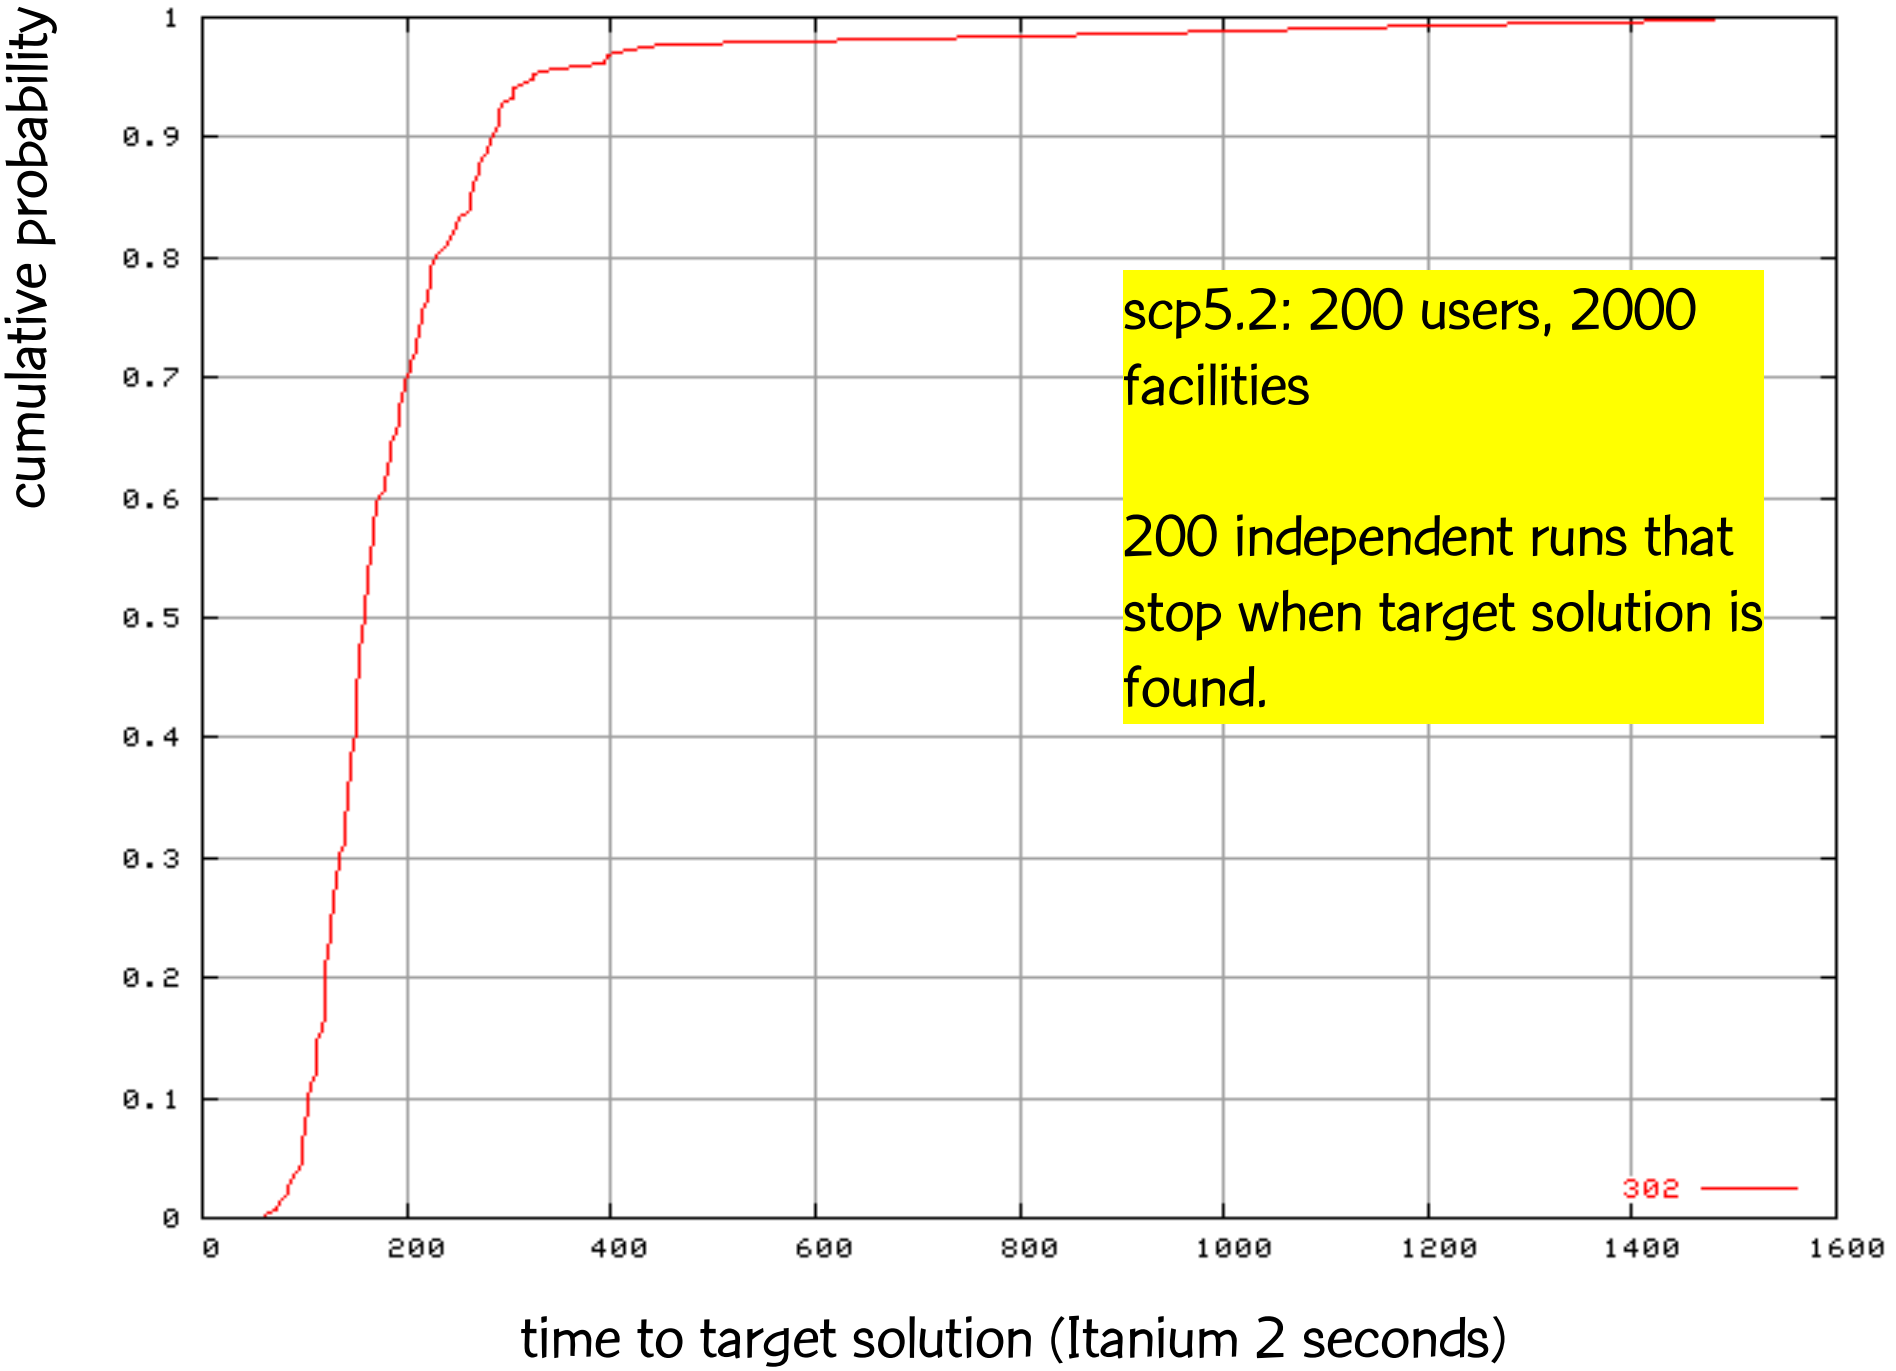
cumulative probability
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
200
400
600
800
1000
1200
1400
1600
302
scp5.2: 200 users, 2000 facilities
200 independent runs that stop when target solution is found.
time to target solution (Itanium 2 seconds)

AT&T

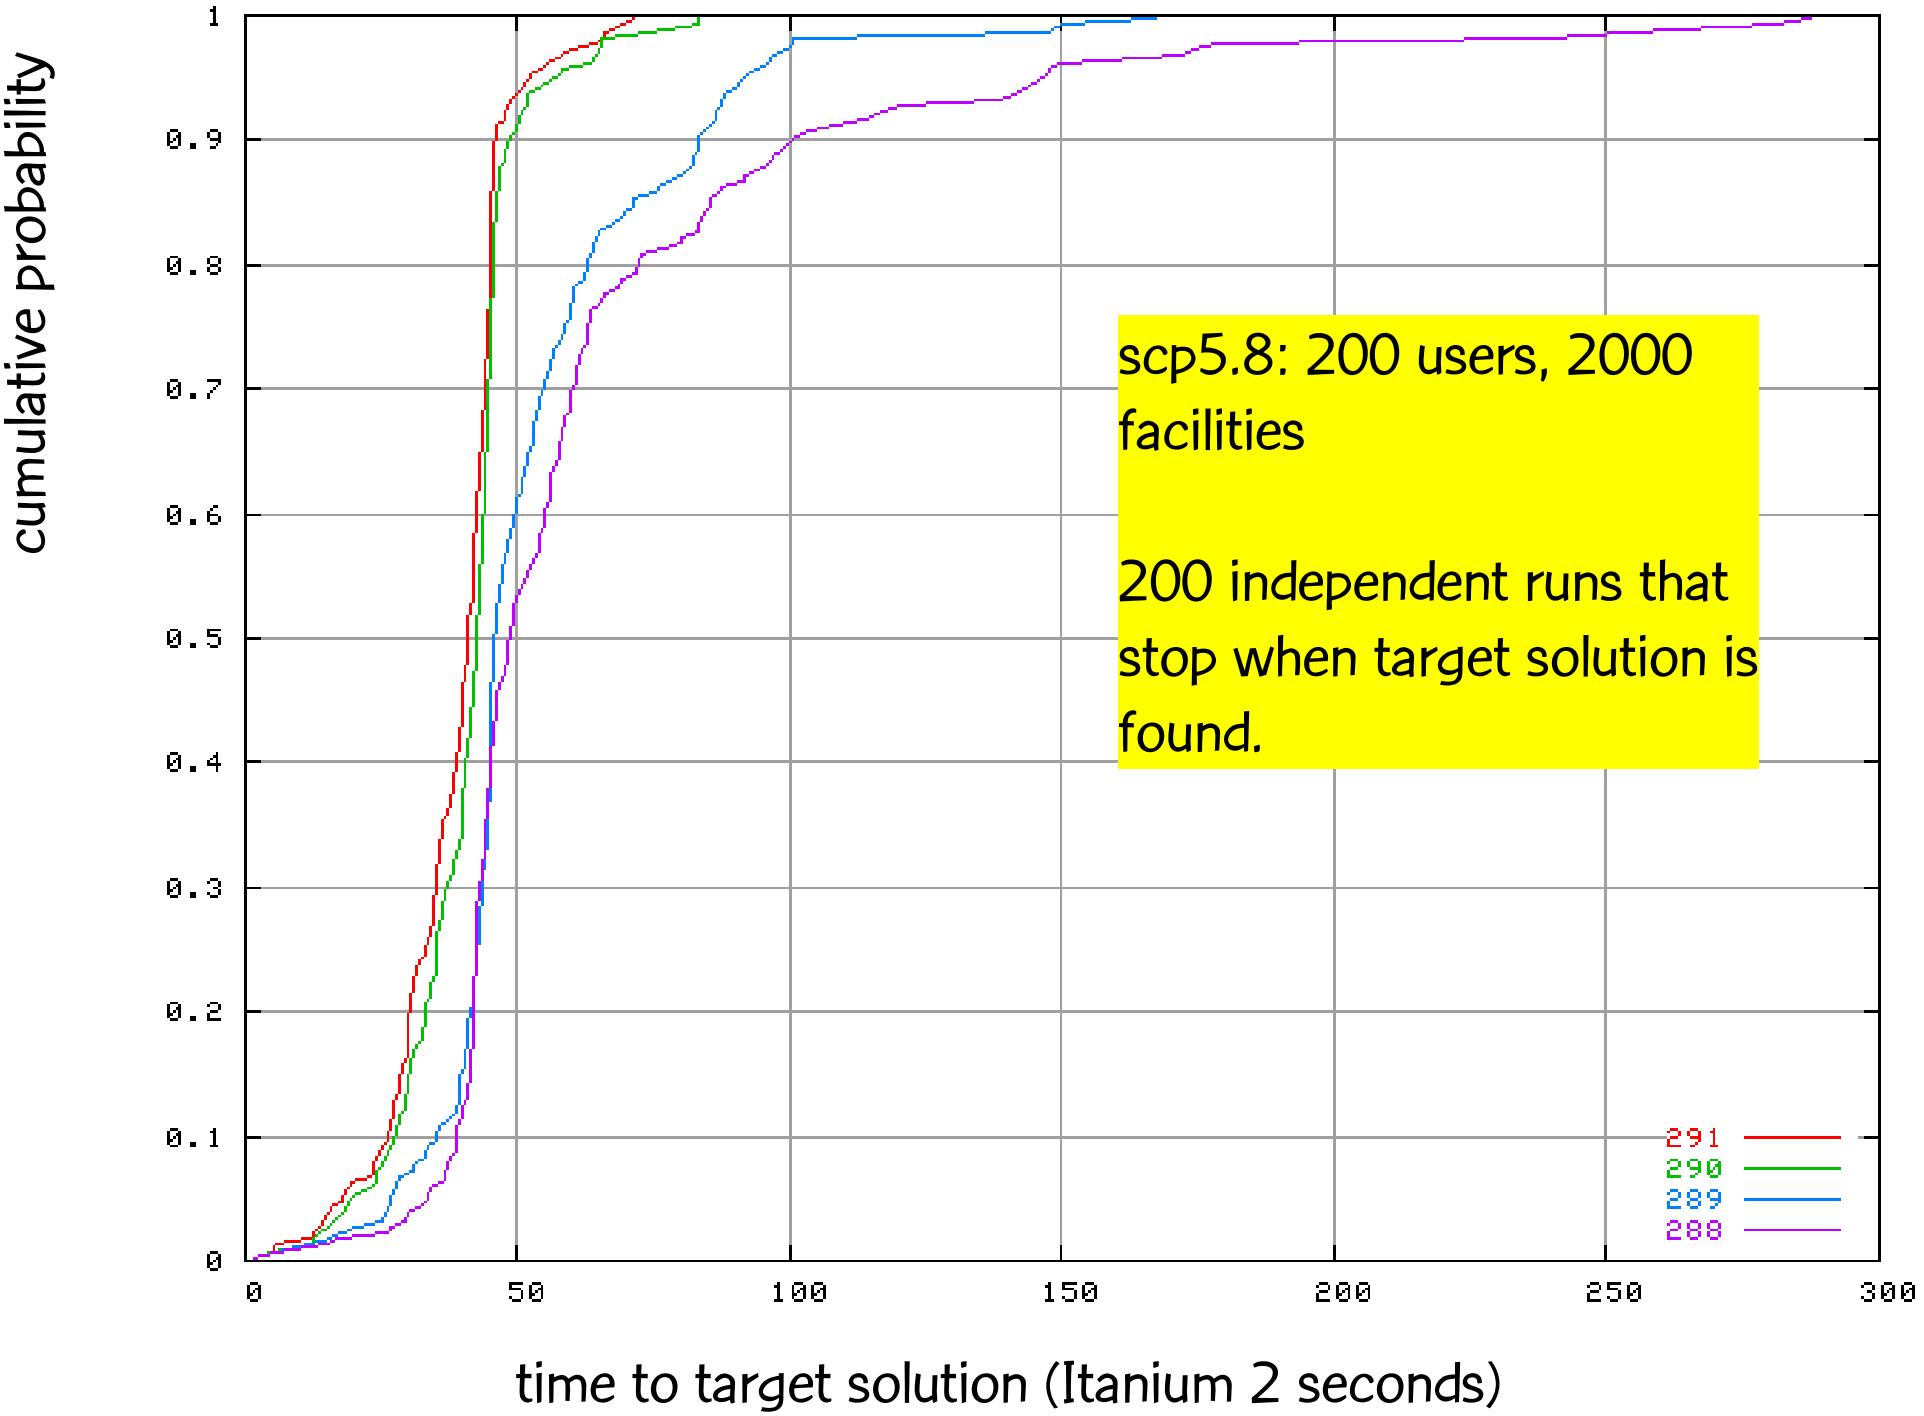
cumulative probability
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
50
100
150
200
250
300
time to target solution (Itanium 2 seconds)
scp5.8: 200 users, 2000 facilities
200 independent runs that stop when target solution is found.
291
290
289
288

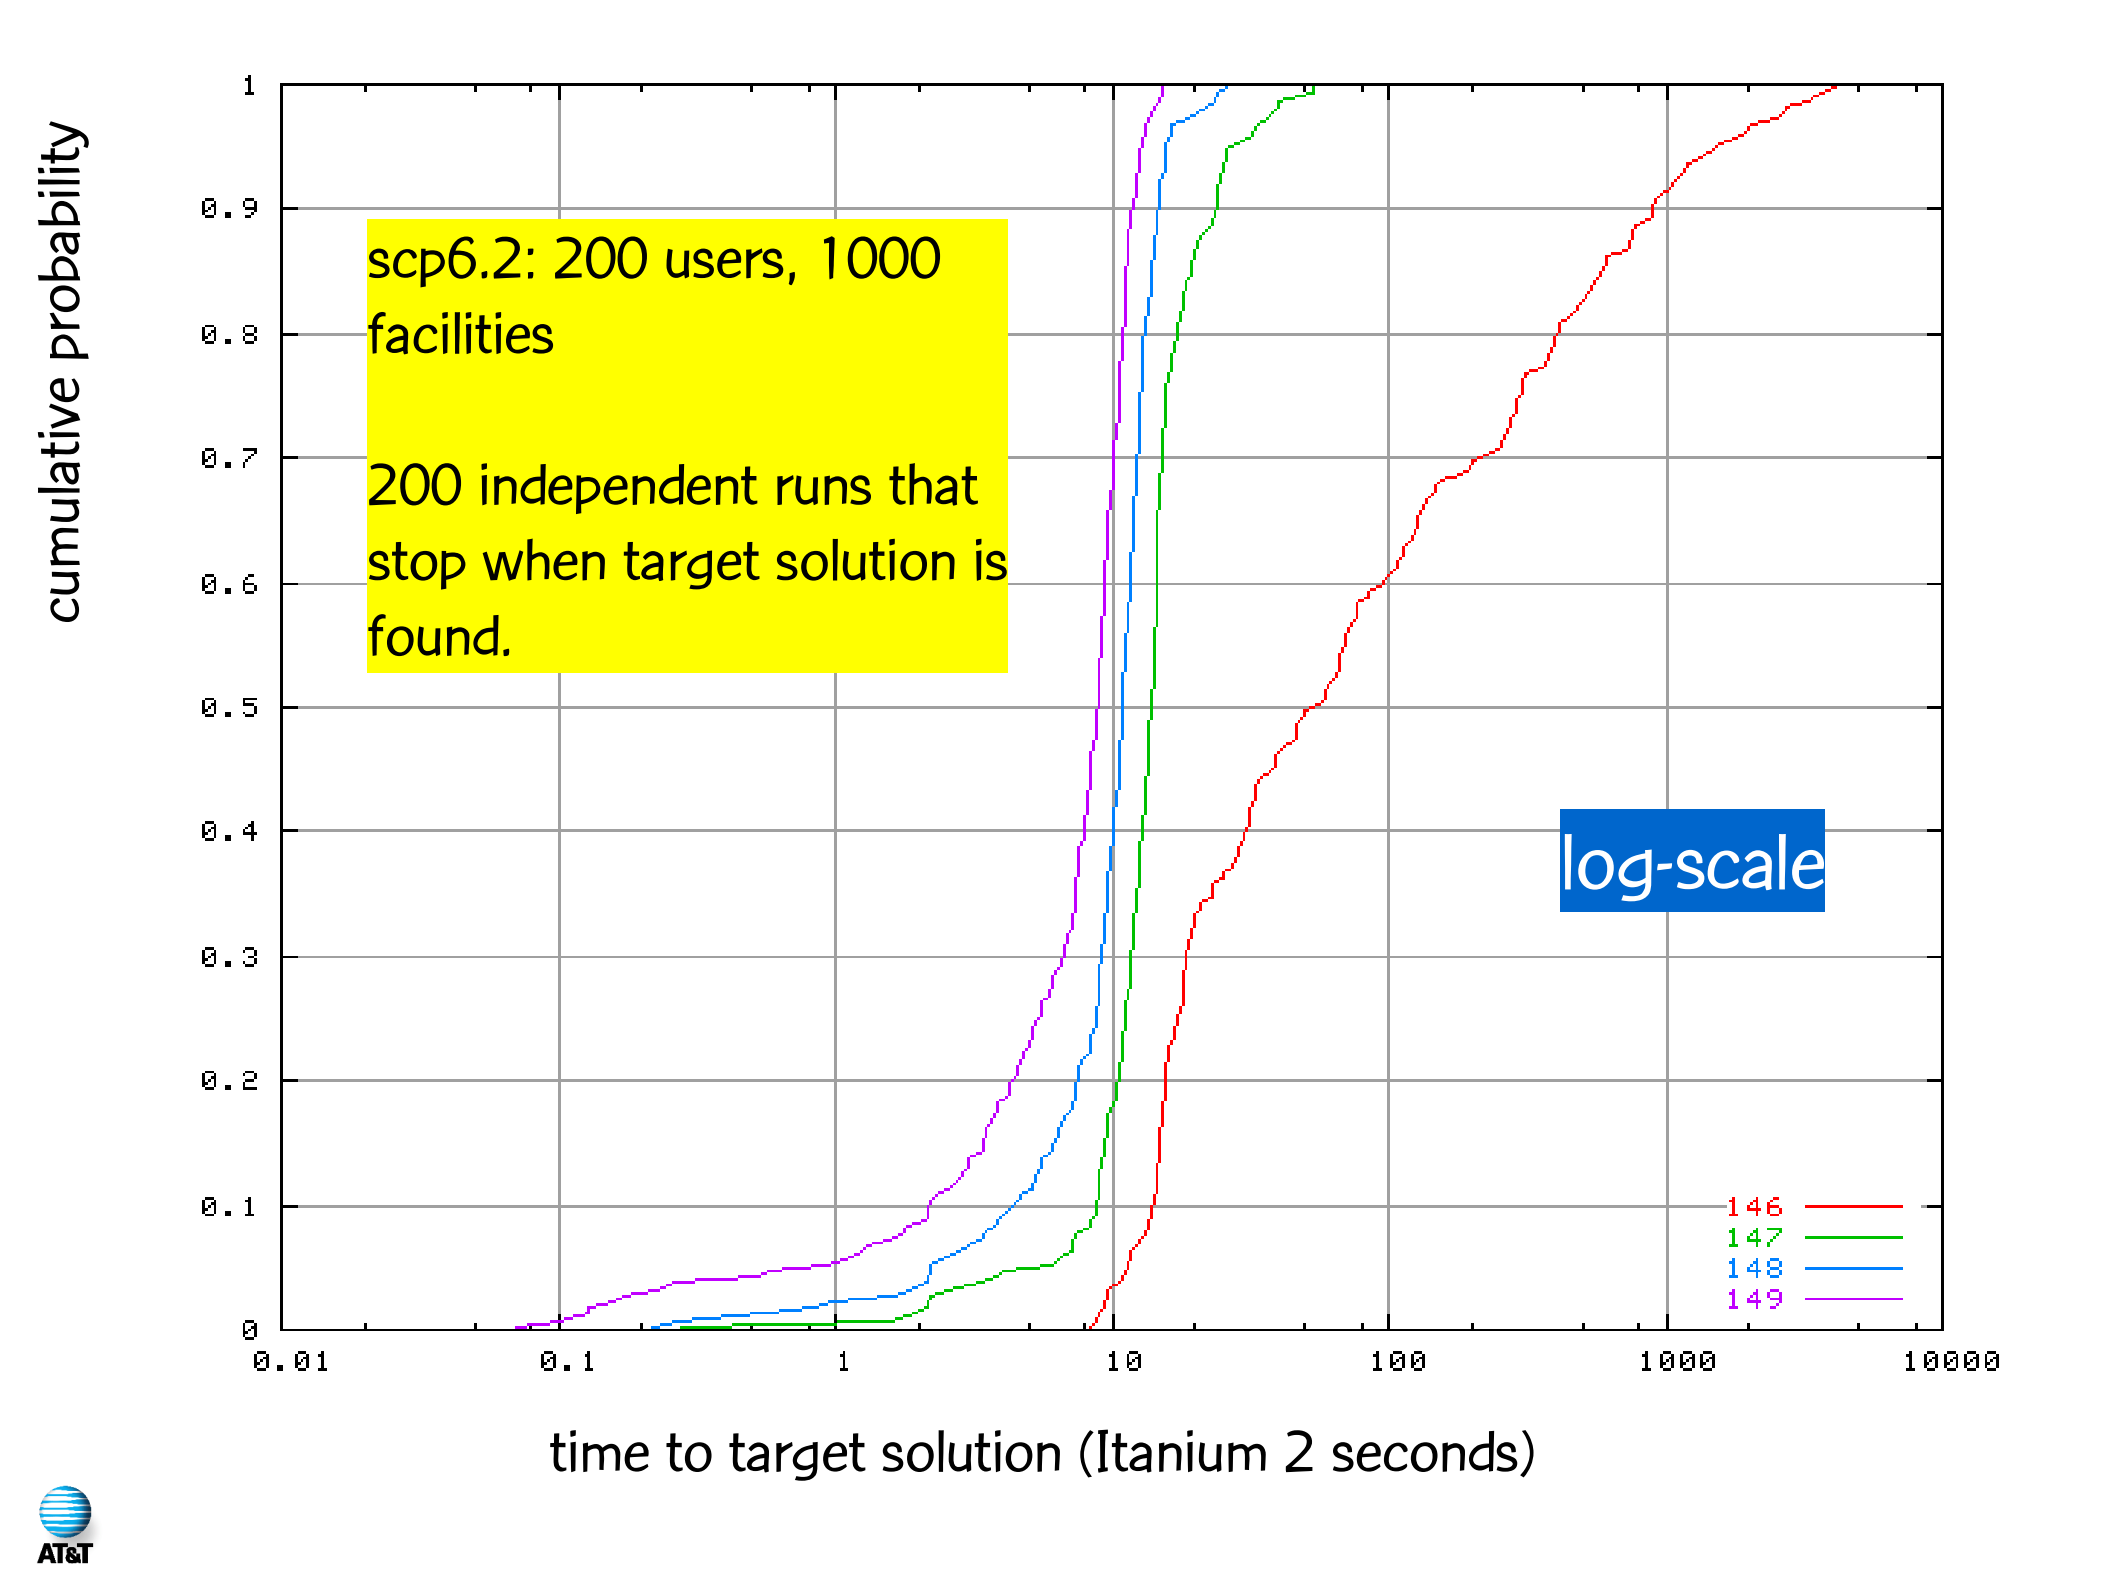
cumulative probability
scp6.2: 200 users, 1000 facilities
200 independent runs that stop when target solution is found.
log-scale
146
147
148
149
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0.01
0.1
1
10
100
1000
10000
time to target solution (Itanium 2 seconds)

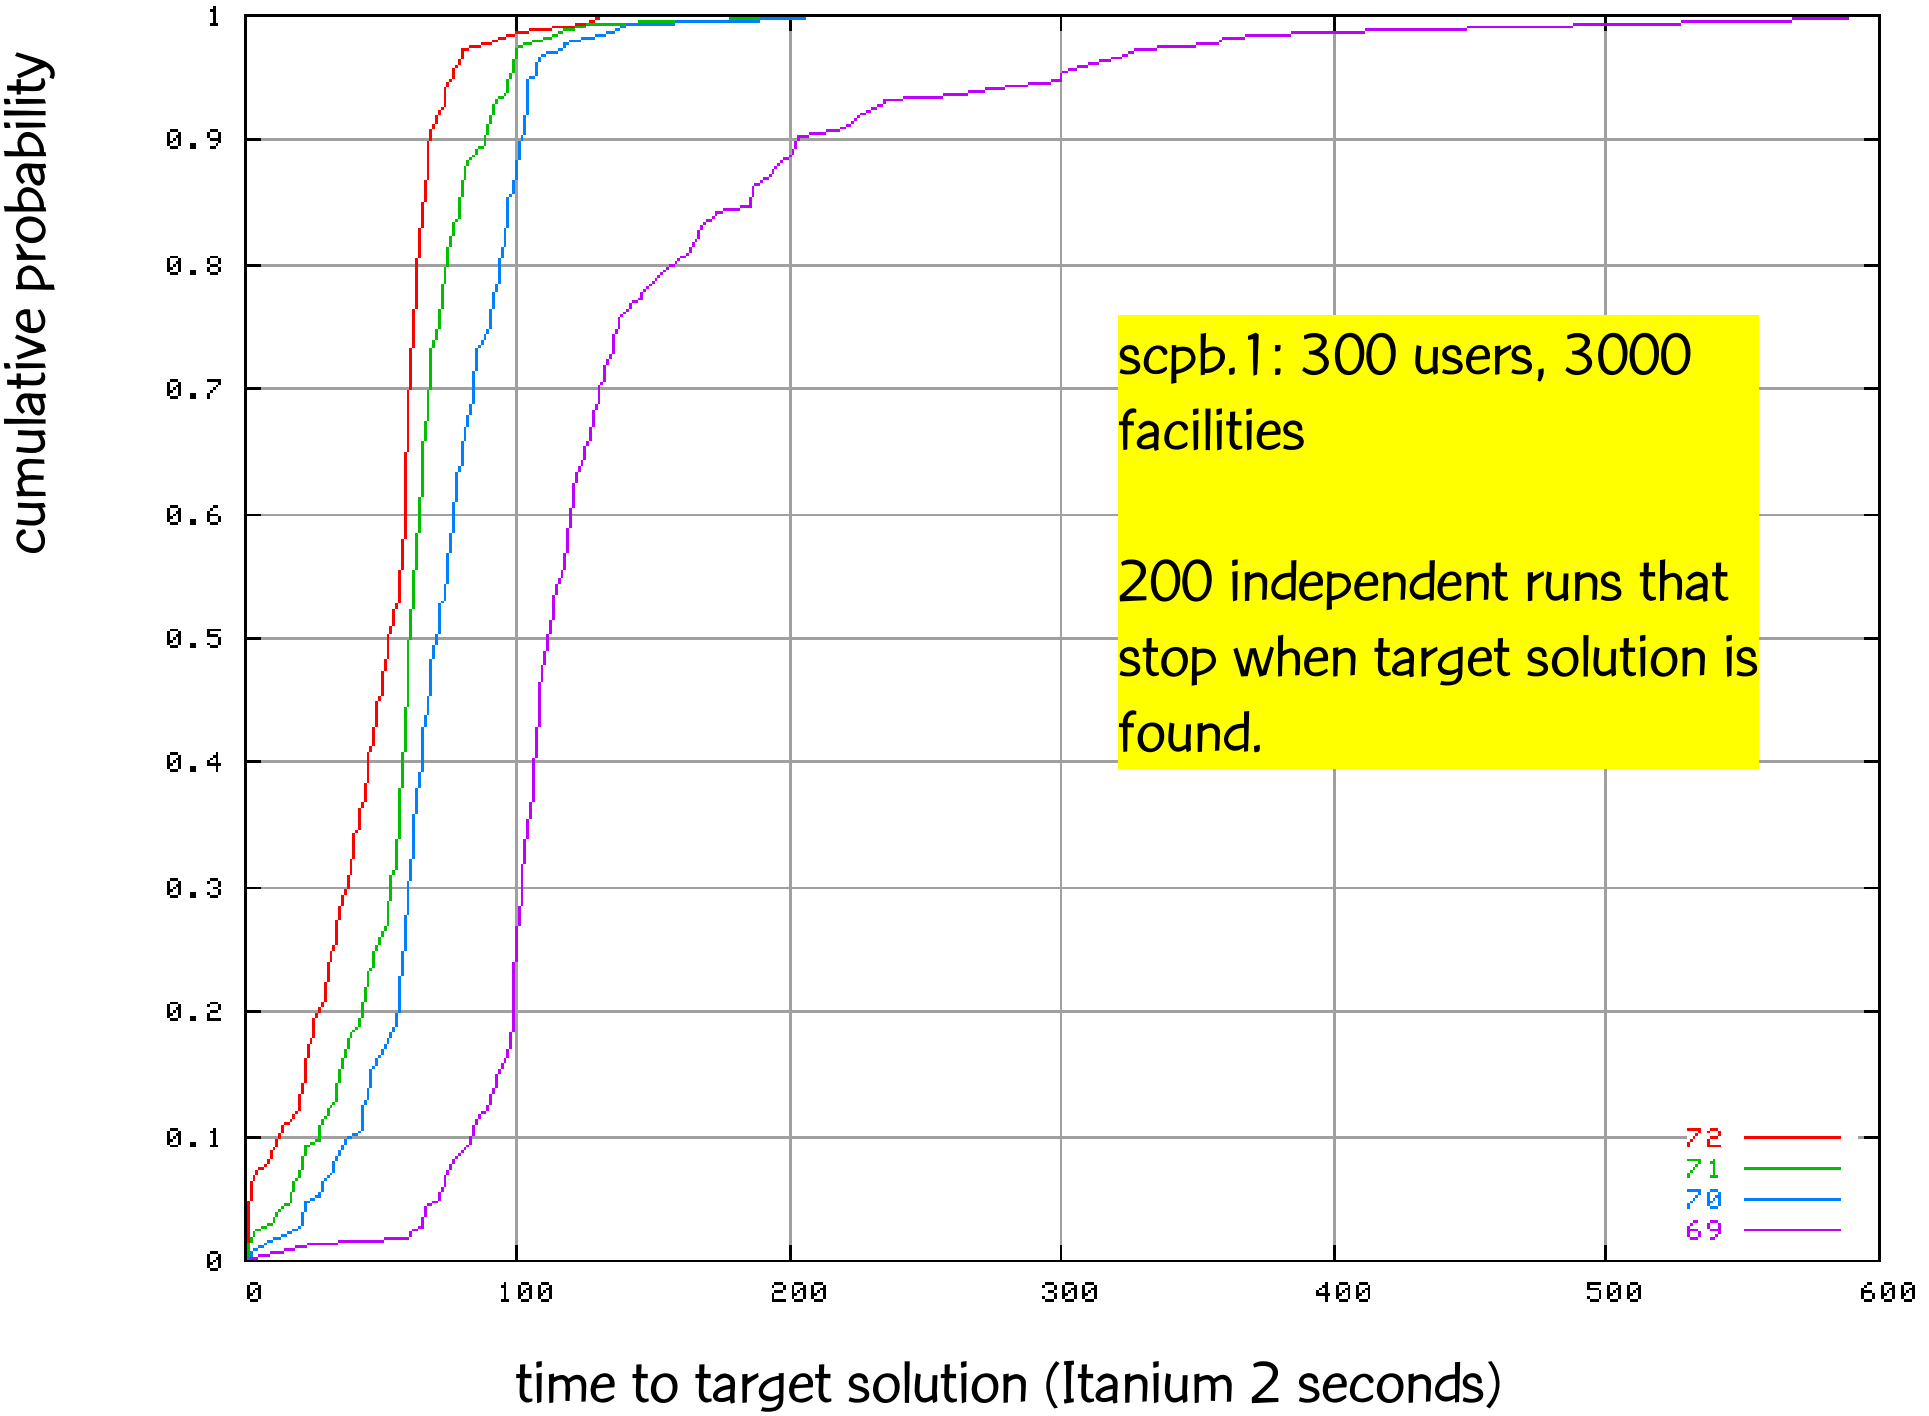
cumulative probability
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
100
200
300
400
500
600
time to target solution (Itanium 2 seconds)
scpb.1: 300 users, 3000 facilities
200 independent runs that stop when target solution is found.
72
71
70
69

AT&T

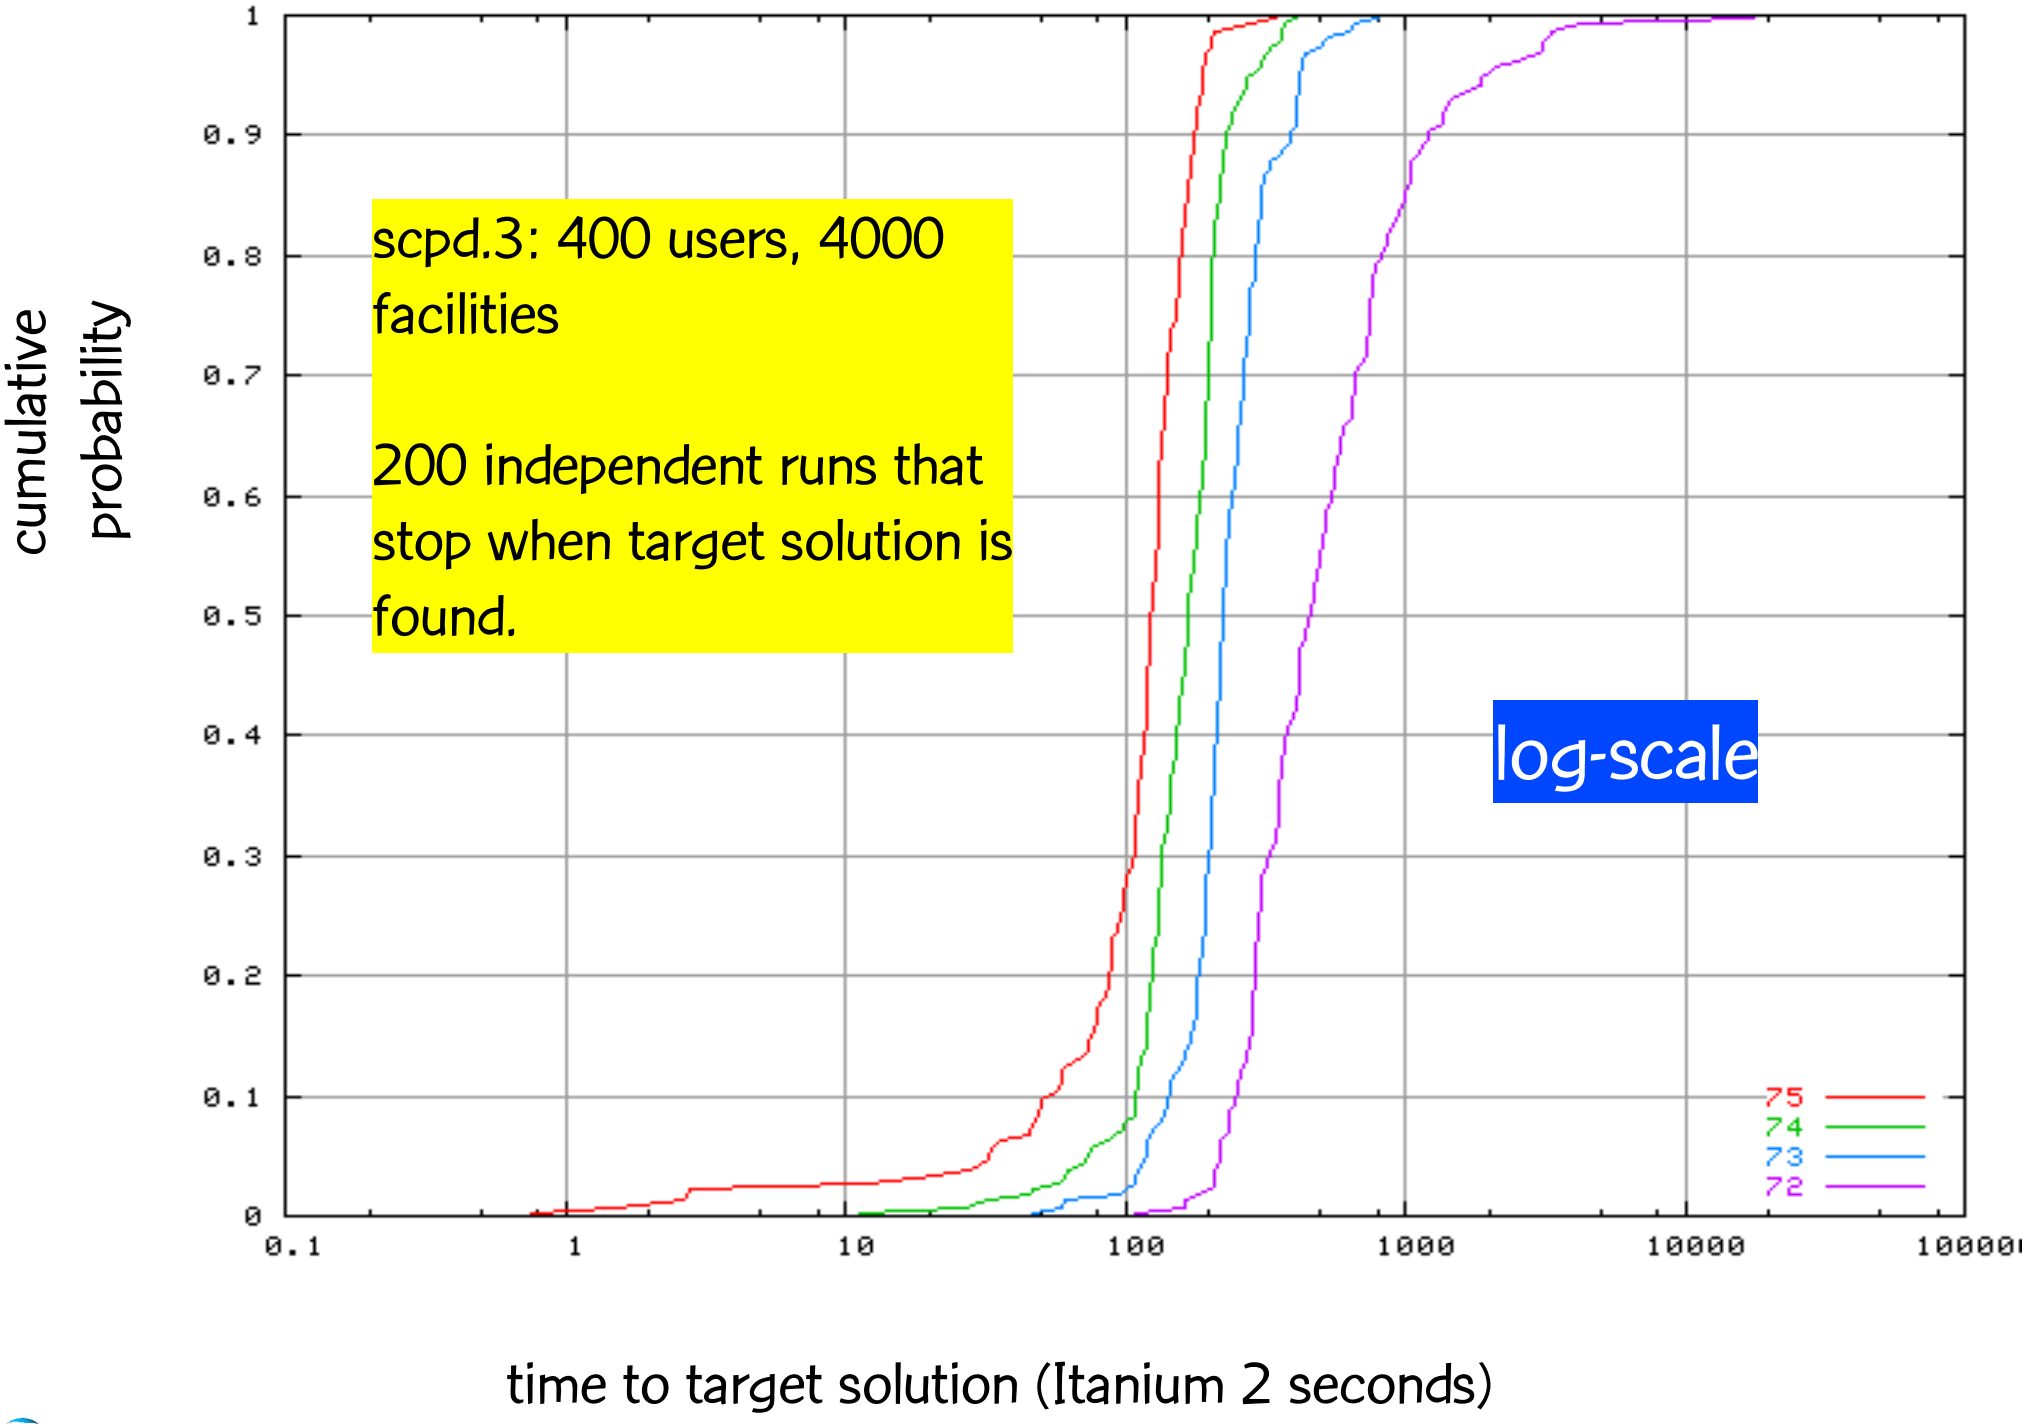
cumulative probability
scpd.3: 400 users, 4000 facilities
200 independent runs that stop when target solution is found.
log-scale
75
74
73
72
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0.1
1
10
100
1000
10000
100000
time to target solution (Itanium 2 seconds)

AT&T

Concluding remarks

- Added to the software `popstar` the ability to solve max covering and set covering problems

  `http://www.research.att.com/~mgcr/popstar`

- A paper with extensive computational experiments is being prepared.

- Weighted max satisfiability problems can also be solved as a facility location problem and will be added to `popstar`.

My coauthors

Diogo V. Andrade
Rutgers University

Renato F. Werneck
Princeton University

AT&T

The End