

ERT, Glue, and Scheduling : present and future

Jeff Templon
NIKHEF



What is ERT?

- * Answers : “If I give you this job right now, how much time will elapse before it starts to execute?”
- * Note : asking the **right question** is important!

Approaches

- * “formula”
- * simulated
- * historical
- * statistical
- * “binary”

Approaches

- * “formula”
- * simulated
- * historical
- * **statistical**
- * “binary”

Statistical

- * No history
- * No "formula"
- * Simple calculations

Copernican Principle

“Your position is not special”

- * Due to J.R. Gott
- * Used in wide variety of cases
 - * Fall of Berlin Wall
 - * Runs of Broadway plays
 - * predictions of “lifetimes” of dictators

Copernicus method and Doomsday theory

[edit]

Gott first thought of his "[Copernicus method](#)" of lifetime estimation in 1969 when stopping at the [Berlin Wall](#) and wondering how long it would stand. Instead of extrapolating a set of developments in world geo-politics ([futurology](#)), Gott used his relative ignorance to his advantage by saying that the [Copernican principle](#) is applicable in cases where nothing is known; unless there was something special about his visit (which he didn't think there was) this gave a 75% chance that he was seeing the wall *after* the first quarter of its life. Based on its age in 1969 (8 years), Gott left the wall with 75% confidence that it wouldn't be there in 1993 $((8/.25) + 1961)$.

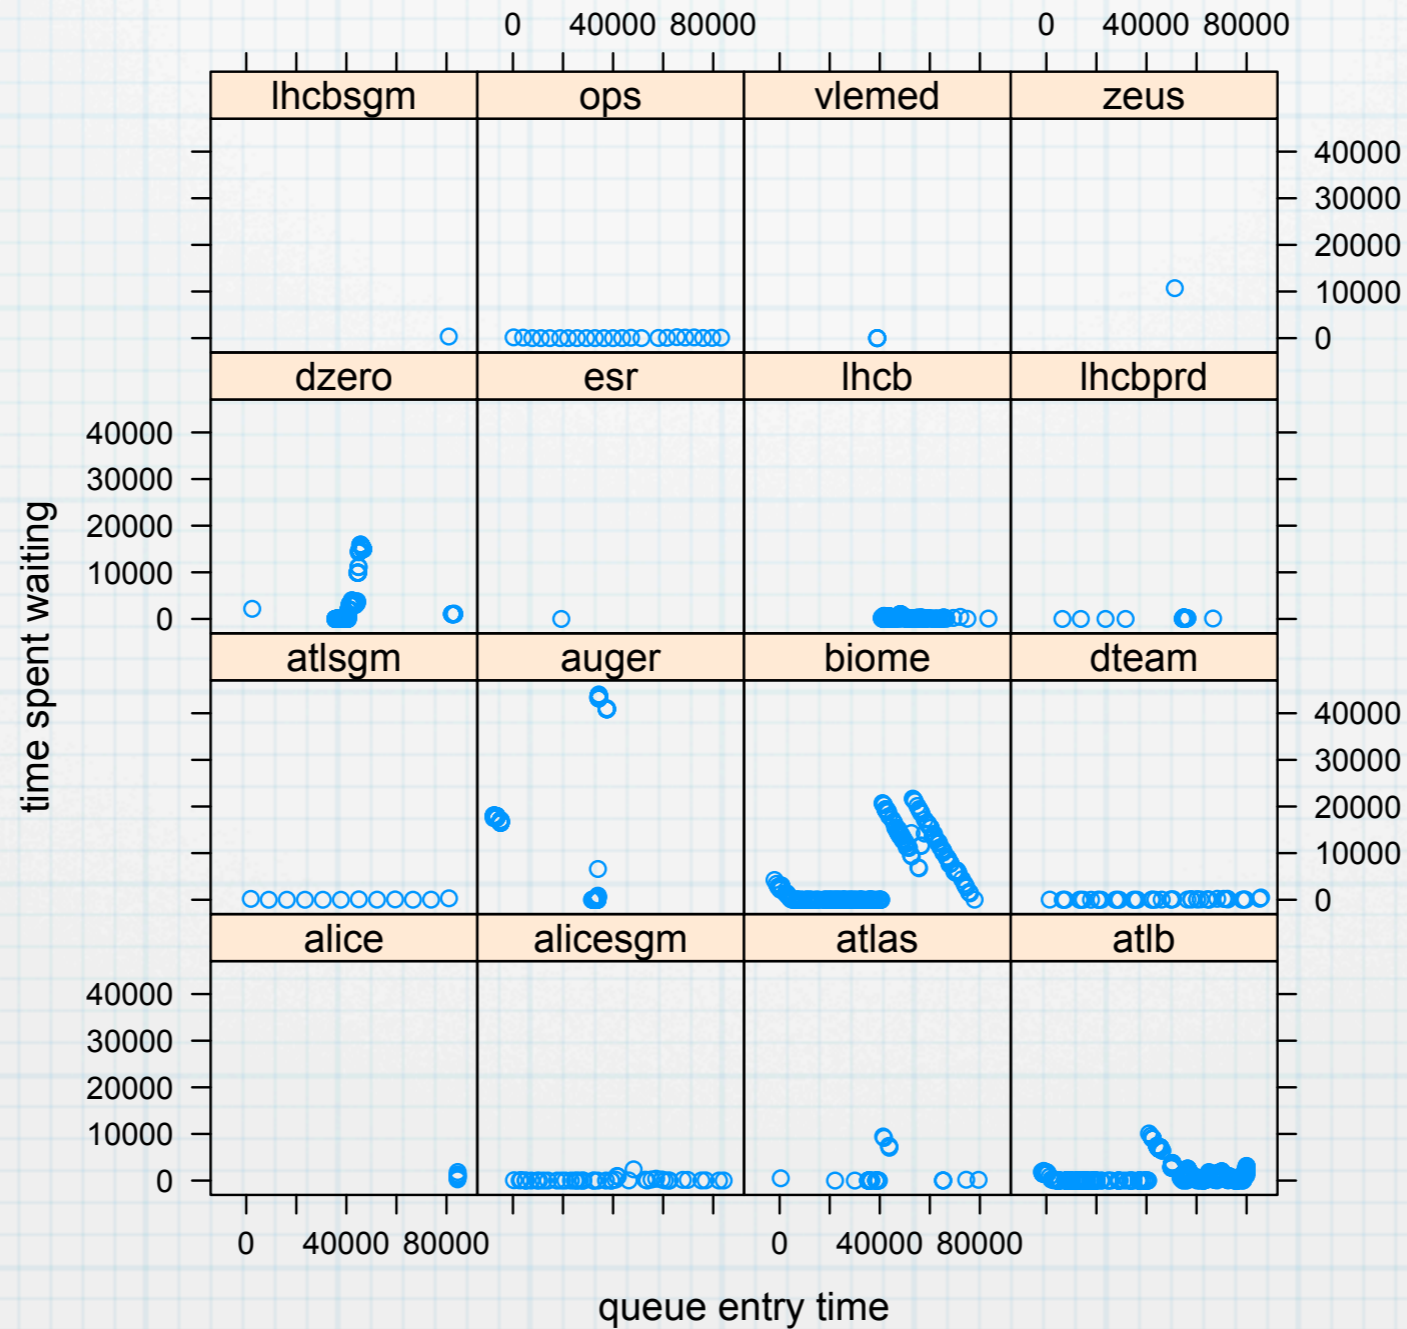
In fact, the wall was brought down in 1989, and 1993 was the year in which Gott applied his "Copernicus method" to the lifetime of the human race. His paper in [Nature](#) was the first to apply the [Copernican principle](#) to the survival of humanity; His original prediction gave 95% confidence that the human race would last for between 5100 and 7.8 million years. ([Brandon Carter's](#) alternative form of the [Doomsday argument](#) was delivered earlier that year, but Gott's derivation was independent.)

He made a major effort subsequently to defend his form of the [Doomsday argument](#) from a variety of [philosophical](#) attacks, and this debate (like the feasibility of [closed time loops](#)) is still ongoing. To popularize the *Copernicus method*, Gott gave [The New Yorker](#) magazine a 95% confidence interval for the closing time of forty-four [Broadway](#) and Off Broadway productions based only on their opening dates. He was more or less 95% correct.


Copernican ERT

- * Your position is not special, so:
 - * if jobs are waiting for your VIO/queue, you will wait as long as they have
 - * if no free CPUs, and you have no jobs waiting, you will get next core (q: "how long since last free core?")
 - * if free CPUs, and you have no jobs waiting, your job will run immediately

24 October 2007




log scale, subset of VOs



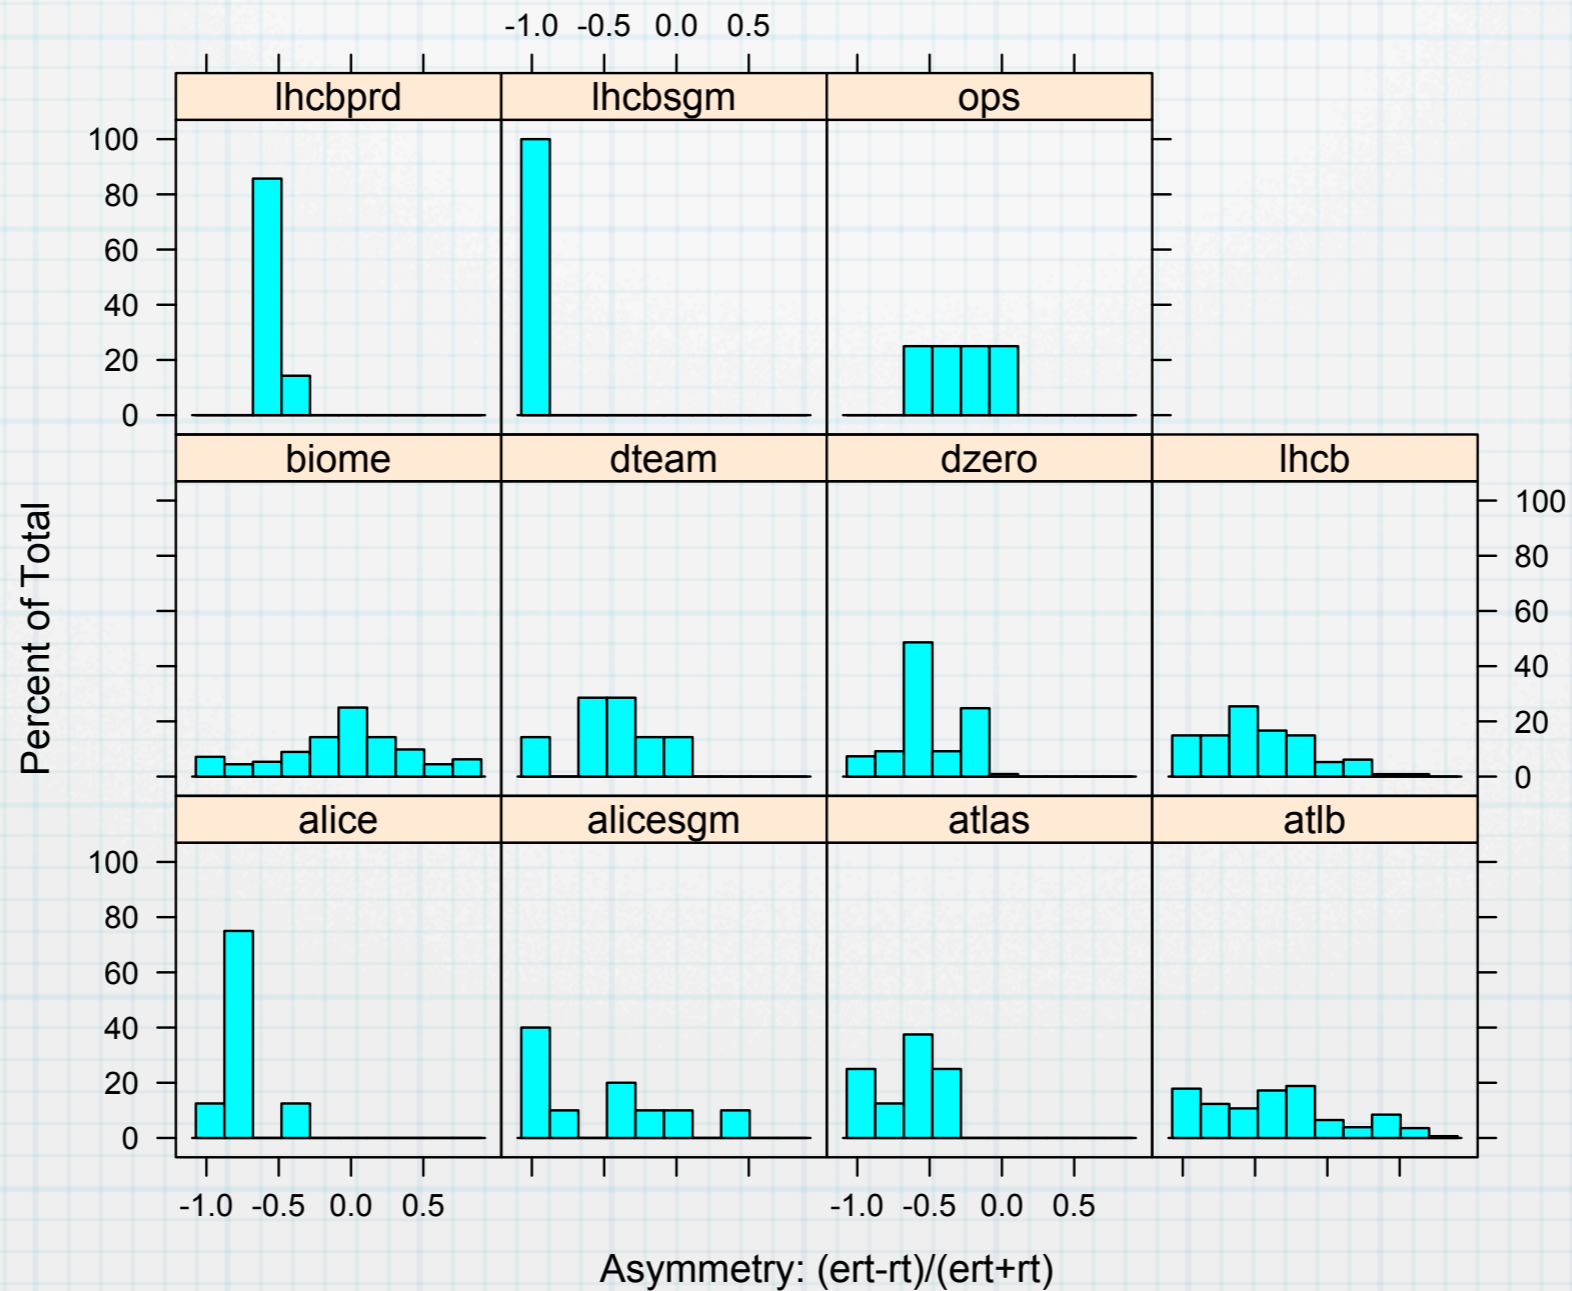
results, wait < 120 sec

Waits < 120 sec, Copernican




results for longer times

Waits > 120 sec, Copernican



ert/rt vs. time

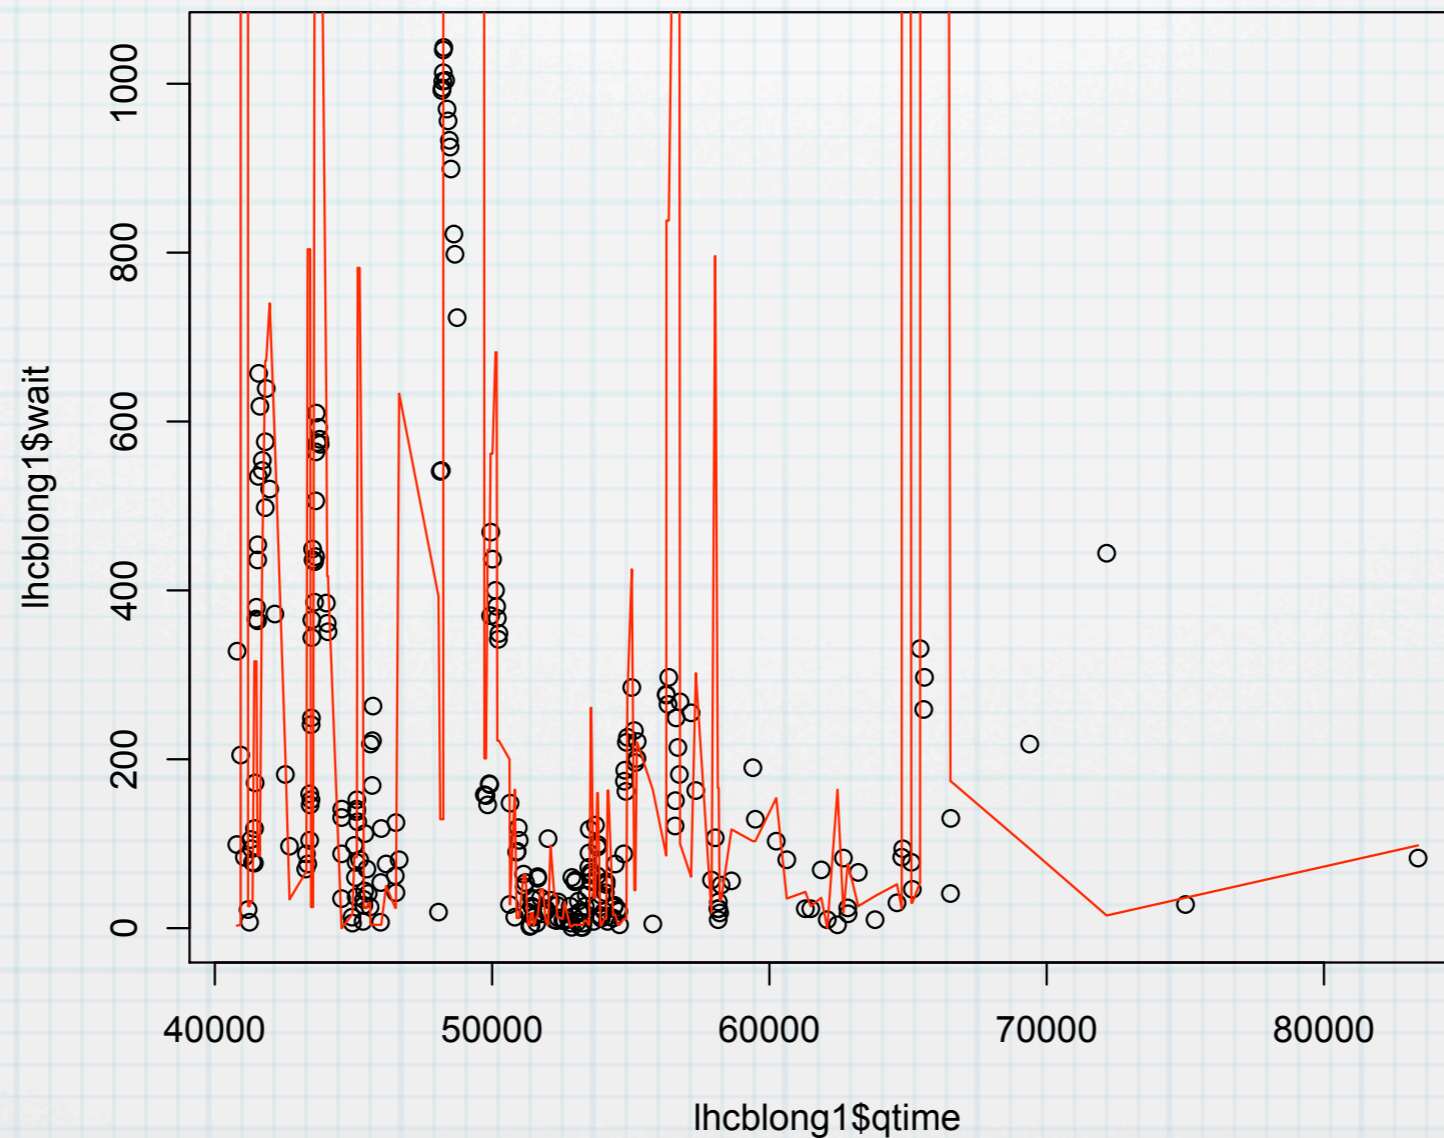


assessment

- * fine for steady state
- * special cases :
 - * queue empties after being full, not bad
 - * jobs suddenly start queuing : **bad**
- * special cases violate assumption, and are also interesting ...
- * It used to be much worse!

Future 1

* there is work on a better algorithm



Development Roadmap

- * fix bugs!
- * new backend (sqlite?) -- current perf limit is searching over job lists
- * new algorithm

Future 2 : missing Glue

- * Q : how many jobs can I give you?
- * now : can only judge by watching info sys ...
- * gradual increase in ERT (as just shown)
- * latencies of info sys (order 10 min)

Possible Good Questions

- * how many jobs could I give you at current ERT?
- * how many slots could I get?
- * how many more slots could I get?
- * note : freeCPUs should always be zero!

Conclusions

- * Current ERT is "OK"
- * does worst at most interesting case
- * the really interesting case* cannot be represented in the info sys!

*for hep-like workflows