Theory Comput Syst (2009) 44: 205–214
DOI 10.1007/s00224-008-9109-y

Die Another Day

Rudolf Fleischer

Published online: 22 March 2008

Abstract The Hydra was a many-headed monster from Greek mythology that would immediately replace a head that was cut off by one or two new heads. It was the second task of Hercules to kill this monster. In an abstract sense, a Hydra can be modeled as a tree where the leaves are the heads, and when a head is cut off some subtrees get duplicated. Different Hydra species differ by which subtress can be duplicated in which multiplicity. Using some deep mathematics, it had been shown that two classes of Hydra species must always die, independent of the order in which heads are cut off. In this paper we identify three properties for a Hydra that are necessary and sufficient to make it immortal or force it to die. We also give a simple combinatorial proof for this classification. Now, if Hercules had known this…

Keywords Hydra battle · Buchholz Hydra · Peano arithmetic · PA · Koenig's lemma · Transfinite induction

1 Introduction

According to Greek mythology, the *Hydra* was a many-headed monster living in a marsh near Lerna [22]. If one head was cut off, one or two new heads grew from the Hydra's body. Nevertheless, in his second task (of twelve, ordered by his cousin Eurystheus) the Greek hero Hercules (a.k.a. Herakles), a son of Zeus, defeated the Hydra, although he did not fully play by the rules: while he was happily hacking away at the heads, his nephew Iolaus burnt the Hydra to prevent new heads from growing [25].

This work was supported by a grant from the National Natural Science Foundation of China (No. 60573025).

R. Fleischer (✉)
Fudan University, Shanghai Key Laboratory of Intelligent Information Processing, Department of Computer Science and Engineering, Shanghai 200433, China
e-mail: rudolf@fudan.edu.cn

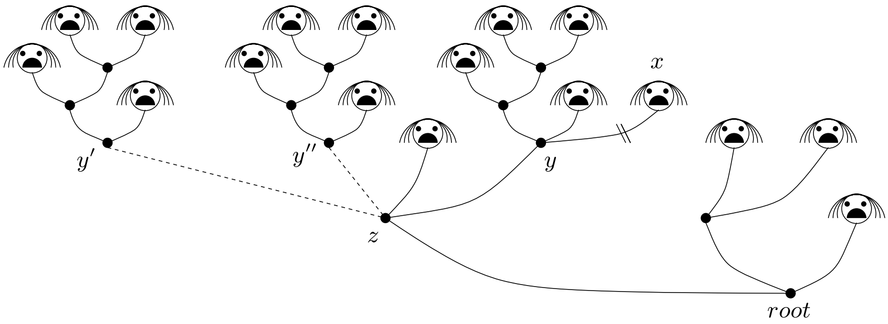

Fig. 1 Cutting off x and growing two copies of T_y^- from z

Kirby and Paris studied this epic fight from a graph theoretic point of view and showed that Hercules might even have won without employing unfair tactics (although, then, he would probably still be fighting today) [20]. They proposed to model the Hydra as a rooted tree where the heads are the leaves (see Fig. 1). The classical Hydra would grow one or two new heads replacing any head that was cut off. Clearly, this Hydra species cannot die (except by fire). Now, Kirby and Paris suggested to study another Hydra species which can duplicate entire subtrees which contained the head that was cut off. At the first glance, one might expect this to be a more powerful Hydra variety, but they proved that this species must eventually die (and this cannot be shown by a simple proof by induction).

To be more precise, they suggested to duplicate subtrees as follows (we call this Hydra species the *i-head Hydra*). For a node v, let T_v denote the subtree rooted at v. For a leaf x, its predecessor y is called the *neck*, and the predecessor's predecessor z is called the *trunk*. The path from the root to y is called the *spine* of x. When x is cut off in the i-th blow, i new subtrees identical to T_y without x, denoted by T_y^-, will grow from the trunk z. Figure 1 shows an example of a second blow. We cut off head x, and the Hydra grows two copies of T_y out of z. If x has no siblings, its neck y becomes a leaf, i.e., a new head, and i new heads (the copies of T_y) grow from the trunk z. If y is the root, no new subtrees grow. If x is the root, i.e., the root is the only node (and head) of the Hydra, cutting off x will *kill* the Hydra.

A Hydra is *doomed* if it will die in a finite number of steps, for any possible sequence of head cuts. Otherwise, it is *immortal*. In a deep mathematical proof based on transfinite induction (if we add the first ordinal number ω, the smallest infinite number, and arbitrary polynomial expressions of ω to the set of natural numbers and define the operation $\omega - 1$ as choosing an arbitrary finite number smaller than ω, then the principle of transfinite induction states that we always reach zero in a finite number of steps when counting down from an arbitrary number) Kirby and Paris showed that the i-head Hydra is doomed. Later, Luccio and Pagli gave an elementary combinatorial proof based on a potential function defined on the nodes for the special case of the 2-*head Hydra* which can only grow two (or any fixed constant number of) subtree copies in each step. They posed as an open problem to find an elementary combinatorial proof for the i-head Hydra [23].

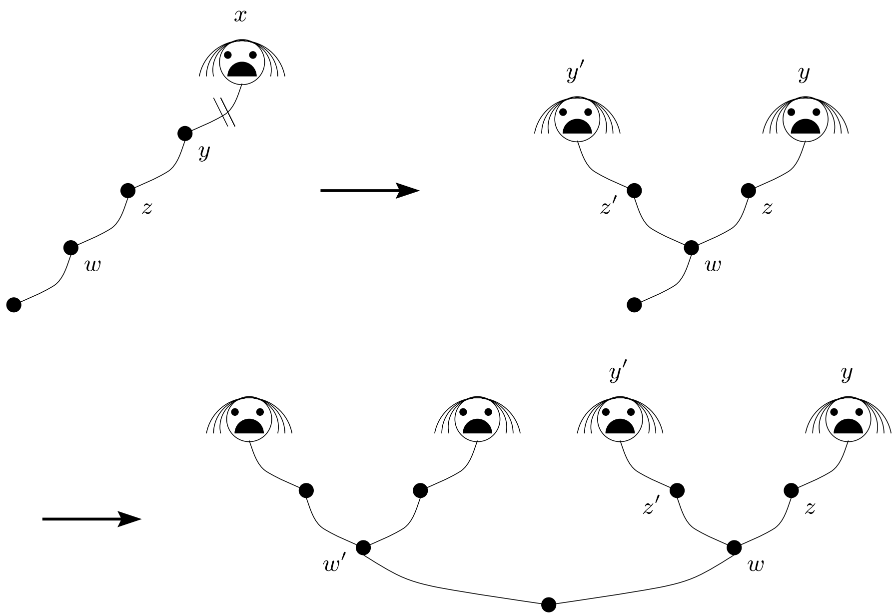

Fig. 2 When we cut head x in a finite Hydra, it can first grow a copy of T_z^- at w and then grow a copy of the new T_w^- at the root

In this paper we give such a proof for the more general class of finite Hydras. As we will see later, the actual number of subtrees grown in each duplication step is not relevant for the fate of a Hydra as long as it is always a finite number, only the locations of the subtrees to be copied do matter.

A *finite Hydra* is a finite tree. If we cut off a head x, it can, in any order, copy an arbitrary but finite number of subtrees according to the following three properties (see Fig. 2 for an example).

(P1) The subtree to be copied is rooted at a node on the spine of x.
(P2) The subtree copy becomes a child of a node on the spine of x.
(P3) The subtree copy is placed at the same level or closer to the root than the original subtree.

In principle, a Hydra may choose not to copy an entire subtree but only part of it (a subtree of the subtree), but obviously it cannot gain anything by doing so, so we may assume w.l.o.g. that a Hydra always copies entire subtrees.

Clearly, the 2-head and i-head Hydras are special cases of the finite Hydra. The number of new subtree copies may either be predetermined (e.g., the 2-head Hydra), given as a function of the structure of the tree or the length of the fight (e.g., the i-head Hydra), or the Hydra may adapt to the cutting sequence, deciding on the number of tree copies each time a head is cut off (this corresponds to choosing an arbitrary finite number in the operation $\omega - 1$ in the transfinite reduction proof by Kirby and Paris).

Generalizing the results by Luccio et al. and Kirby et al., we show that any finite Hydra is doomed.

Theorem 1 (Hydra Theorem) *Any finite Hydra is doomed.*

We will give a simple combinatorial proof of the Hydra Theorem in Sect. 3, after shortly reviewing the mathematical history of this problem in Sect. 2. In Sect. 4 we will see that relaxing any one of the three properties (P1)–(P3) can make a Hydra immortal. In Sect. 5 we shortly discuss worms [15], a one-dimensionally restricted Hydra, and the Buchholz Hydra [3], a truly two-dimensional generalization of the finite Hydra.

2 The History of the Hydra Battle

Kurt Tucholsky once noted a speaker should "always start with ancient Rome and mention the historical background of the matter" [29]. We already discussed the pre-Roman story of Hercules and the Hydra, so we may jump directly to the past century (there was not much happening in between; at least, nothing related to this paper).

There is a branch of mathematical logic that is concerned with the relative power of the various axioms of mathematics [11] (the Axiom of Choice (AC), the axioms of Peano Arithmetic (PA), etc.). After Gentzen had shown the consistency of PA [13] (see also [28]) in 1936, Goodstein gave in 1944 a recursive definition of the so-called Goodstein sequence and showed that in PA it cannot be proven to terminate [14] (see also Cichon [5]). Basically, this means there is no classical proof by induction for this theorem. A problem of a similar flavour is the famous $3x + 1$ conjecture (where termination has not been proven yet), also known as Collatz problem, Syracuse problem, Kakutani's problem, Hasse's algorithm, and Ulam's problem [21].

Much later, in 1982, Kirby and Paris gave an alternative proof for the termination of the Goodstein sequence [20] and introduced the Hydra battle as another example to demonstrate their new technique from [24]. They showed that i-head Hydras are doomed and this cannot be shown in PA, even if Hercules is required to always cut off the rightmost head of the Hydra (assuming a natural ordering of subtrees by decreasing size). In this case, even if the Hydra has only height two, the length of the battle is not primitive recursive. Another proof was given by Carlucci via reduction from Gentzen's Reduction Strategy [4].

The Kirby-Paris paper led to a flurry of research on the Hydra and related problems. It was quickly observed that their results actually hold for the more general case of *arbitrary-head Hydras* (which decide in every step how many subtree copies are grown). Generalizations to growing copies at arbitrary nodes on the spine (as in our finite Hydra) are implicit in Jouannaud's survey paper on proofs and computation [18]. Dershowitz and Moser recently gave a survey on the Hydra battle formulated as a rewriting system [9].

Since the Kirby-Paris Hydra can only grow in width, the question arose whether there is a generalization of these results to height-growing Hydras. In 1987, Buchholz answered this in the positive with a doomed Hydra species that can grow in width and height [3] (basically, the growth in each dimension is bounded by a Hydra battle). This Hydra is now known as the *Buchholz Hydra*.

Hamano and Okada went the other direction and restricted Hydras to truly one-dimensional objects, so-called worms [15] (see also Beklemishev [2]). Although being recursively defined and of reasonably bounded length, they cannot, in PA, be proven to terminate. Hydras also came to fame in the Scientific American when Gardner [12] discussed the Hydra battle and Smullyan's Urn Game [26], kind of a simplified version of the arbitrary-head Hydra.

In 2000, Luccio and Pagli discovered the combinatorial beauty of the Hydra battle [23] and asked whether there is a simple combinatorial proof that the i-head Hydra is doomed. Luccio's presentation of the problem at FUN 2001 let to lively discussions among the conference participants (futily trying to find proofs) and might be considered the birth hour of the present paper. Here, we give a simple combinatorial proof that finite Hydras are doomed. We use Koenig's Lemma, which pops up here and there in the computer science literature, mainly in proofs in logics and formal languages. Although it is deceivingly simple to state and prove, it is actually a powerful theorem outside of PA, equivalent to AC [6, 17, 19], so our simple proof does not contradict the Kirby-Paris result. It is interesting to note that the lemma does not generalize to the next higher level of infinity. An uncountable tree where each level is countable does not necessarily have an uncountable path; an example of such a tree are the Aronszajn trees [10].

Lemma 2 (Koenig's Lemma) *A tree is finite if and only if every node has finite degree and every simple path from the root is finite.*

Proof If one node has infinite degree or there is an infinite simple path, the tree is infinite. On the other hand, if a tree where all nodes have finite degree is infinite, then one of the subtrees of the root must be infinite. If we follow the edge to that subtree and iterate, we can construct an infinite simple path. □

Daly also used Koenig's Lemma to give simple proofs for Smullyan's Urn Game and the Goodstein sequence [7]. Weiermann recently showed the exact threshold of the transition between PA-provability (the 2-head Hydra) and non-PA-provability (the i-head Hydra) [31]. Readers interested in reading the mathematical papers might first want to read some good introduction into the theory of ordinals (for example, Avigad [1] or Dershowitz [8]).

3 Proof of the Hydra Theorem

If not all finite Hydras are doomed, there must be an immortal Hydra H of smallest height. We call the subtrees of the root of H *subhydras*. If we cut a head in a subhydra, H may choose to create a finite number of copies of the subhydra, which we also call subhydras.

Let ρ be an infinite hacking sequence that does not kill H. Now we define the *hacking tree* S. Initially, S consists of a root with one leaf for each subhydra of H.

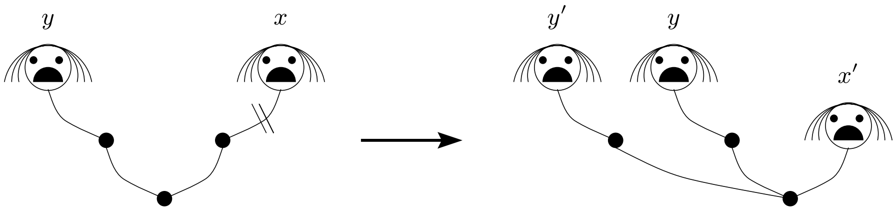

Fig. 3 A Hydra that can copy a subtree not containing the head that was cut off, violating (P1). If we cut off x and the Hydra copies the subtree containing y, the new Hydra is a supertree of the original one, i.e., it is immortal

Each time we cut a head in a subhydra, the corresponding leaf of S will become an internal node with $k+1$ new children leaves, where $k \geq 0$ is the number of copies we create of the subhydra. One of the new leaves corresponds to the subhydra that was copied, the other leaves correspond to the copies. So after each step, each leaf of S corresponds to one of the current subhydras of H, and siblings in S represent identical copies of the same subhydra (here we use (P1), that a subtree can only be copied if it contained the head that was cut off).

Since ρ is infinite, S is infinite. Each node has finite degree (here we use that we only create a finite number of subtree copies in each step), thus there must exist an infinite simple path p in S by Koenig's Lemma, corresponding to an infinite subsequence σ of ρ. For each node v on p, we may w.l.o.g. assume that the successor w (the child of v in S) is actually the original subhydra that was copied in this cutting step (here we use (P3) in a subtle way; our proof by contradiction is actually a proof by induction on the height of doomed Hydras, where the induction step cuts off the root. Therefore, we cannot allow that subtrees get copied higher up in the tree, because otherwise in the inductive step we could have new subhydras appearing out of the blue sky). Thus, σ is an infinite cutting sequence in one of the initial subhydras G of H. Note that (P2) implies that none of the cuts in $\rho - \sigma$ can affect G. So we have found an immortal Hydra G of smaller height than H, a contradiction.

4 Immortal Hydras

In this section we show that the properties (P1)–(P3) cannot be weakened. Violation of any one of them can make a Hydra immortal. Note that the Hydras violating (P1) and (P3) are immortal in a strong sense: they can survive *any* cutting sequence.

We start with a Hydra that is allowed to grow copies of subtrees that did not contain the head that was cut off. But then the Hydra in Fig. 3 can in each step grow to a Hydra that contains the original one as a sub-Hydra, i.e., the Hydra is immortal.

Next we violate property (P2). The Hydra in Fig. 4 is allowed to grow copies of the subtree (containing the head that was cut off) anywhere in the tree. But even if it grows the subtree on a valid level (i.e., not higher than its original level), there is a sequence of two cuts that allows the Hydra to become larger than before, i.e., Hercules may lose this fight.

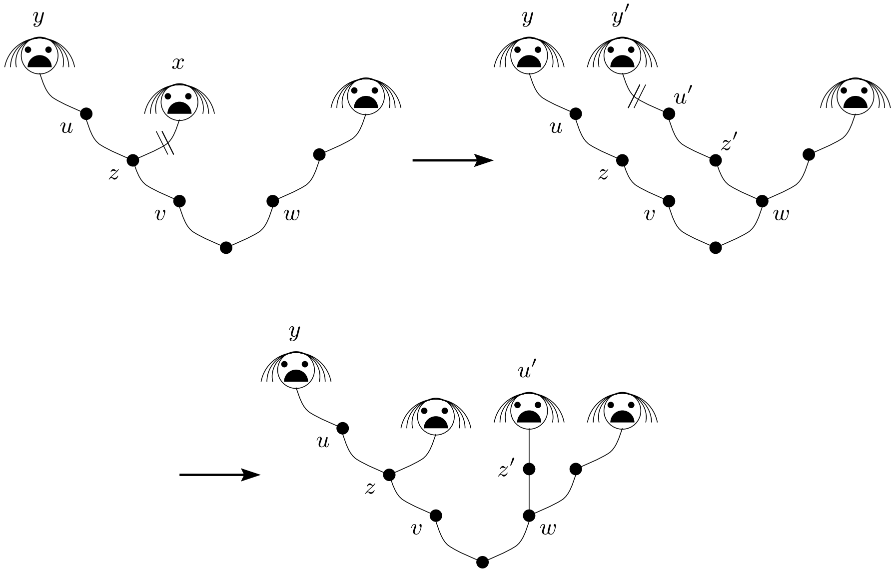

Fig. 4 A Hydra that can copy a subtree to nodes not on the spine of the cut head, violating (P2). If we cut off head x, the Hydra can grow a copy of T_v^- at node w (which is at the same level as v). If we now cut off the new head y', u' will become a head and we can grow a copy of this head at z. This Hydra is now a supertree of the original one, so it is immortal

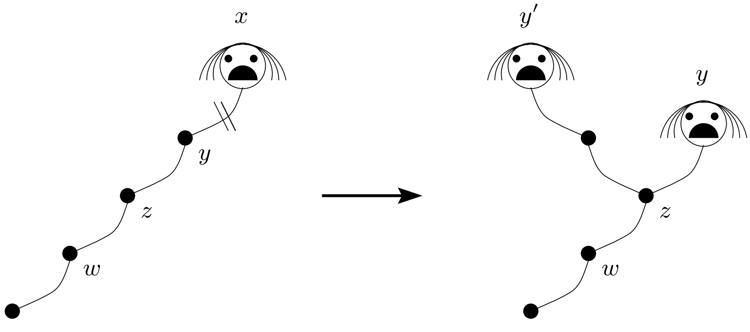

Fig. 5 A Hydra that can place a copy of a subtree higher up on the spine, violating (P3). This is the same Hydra as in Fig. 2, but this time we grow the copy of T_z not at w but at z. The new Hydra is a supertree of the original one and thus immortal

Finally we violate property (P3). The Hydra in Fig. 5 is allowed to grow copies of the subtree at a higher level. But then the resulting Hydra is larger than before, i.e., it is immortal. In the next section we will define the Buchholz Hydra which violates (P3) but is still doomed because the height growth is restricted to be equivalent to an elementary Hydra battle.

5 Worms and the Buchholz Hydra

Hamano and Okada defined a worm as a one-dimensional version of the two-dimensional Buchholz Hydra and showed it to be equivalent to the i-head Hydra [15]. Here we give the self-contained definition by Beklemishev [2] (using a slightly modified notation). A *worm* of length $\ell + 1$ is a finite sequence of natural numbers $w = (f(0), f(1), \dots, f(\ell))$. $f(\ell)$ is called the *head* of the worm. The worm battle is defined by the sequence of worms $w_0 = w$ and $w_{n+1} = \mathit{next}(w_n, n+1)$, defined as follows. Let w be a worm with head $f(\ell)$ and $m \geq 1$ an integer.

1. If $f(\ell) = 0$, then $\mathit{next}(w, m) = (f(0), \dots, f(\ell - 1))$. That is, in this case we cut off the head of the worm.
2. If $f(\ell) > 0$, let $k = \max_{i<\ell}\{f(i) < f(\ell)\}$, i.e., the index of a worm component closest to the head of smaller value. The worm w, with the head decreased by one, is then the concatenation of two parts, the *good* part $r = (f(0), \dots, f(k))$ (which can be empty), and the *bad* part $s = (f(k+1), \dots, f(\ell-1), f(\ell)-1)$. We define $\mathit{next}(w, m) = r * \underbrace{s * s * \cdots * s}_{m+1 \text{ times}}$.

Thus, w_{n+1} is either w_n minus its head, if the head had value 0, or the head is decreased by one and we extend w_n by adding $n + 1$ copies of the bad part of w_n. Note that w_n is defined by a primitive recursive function, and its length is bounded by $|w_n| \leq (n + 2)! \cdot |w_0|$. Still, proving that any worm must eventually die cannot be shown in PA, the Peano Arithmetic (actually, it is equivalent to the i-head Hydra battle).

Buchholz generalized the i-head Hydra to a species that can also grow in height [3] by relaxing (P3). He allowed subtree copies to be placed higher up in the tree, but only a bounded number of times. Since this bound is essentially given by the length of a i-head Hydra battle, this produces a Hydra that grows as high as it grows wide. To be more precise, a Buchholz Hydra is a finite labeled tree H with the following properties:

1. The root has label $+$.
2. Any other node of H is labeled by some ordinal $v \leq \omega$.
3. All nodes immediately above the root of H have label 0 (zero).

If we cut off head x, H will choose an arbitrary number $n \in \mathbb{N}$ and transform itself into a new Hydra $H(x, n)$ as follows. Let y denote that node of H which is immediately below x (the neck), and let H^- denote that part of H which remains after x has been cut off. The definition of $H(x, n)$ depends on the label of x.

Case 1: label$(x) = 0$. If y is the root of H, we set $H(x, n) = H^-$. Otherwise, $H(x, n)$ results from H^- by growing n copies of H^-_y from the node immediately below y, i.e., the trunk. Here, H^-_y denotes the subtree of H^- rooted at y.

Case 2: label$(x) = u + 1$. Let z be the first node below x with label $v \leq u$. Let G be the tree which results from the subtree H_z by changing the label of z to u and the label of x to 0. $H(x, n)$ is obtained from H by replacing x by G. In this case, $H(x, n)$ does not depend on n.

Case 3: label$(x) = \omega$. $H(x, n)$ is obtained from H simply by changing the label of x: ω is replaced by $n + 1$.

Note that a Buchholz Hydra with all labels equal to zero is just the classical i-head Hydra. Buchholz showed that every Buchholz Hydra can be killed by repeatedly cutting off the rightmost head, but this cannot be shown in PA (it can be shown in (Π^1_1−CA)+BI, if you really want to know). Later, Hamano and Okada showed that the Buchholz Hydra is actually doomed with any cutting sequence [16]. Wainer further generalized the underlying mathematics of structured tree-ordinals [30], but it is not clear (to us) what this implies for the Hydra battle.

6 Conclusions and Open Problems

We have characterized the elementary Hydras that are doomed or immortal by identifying three properties (P1)–(P3) that are necessary and sufficient for a Hydra to be doomed. We could not find an example of a Hydra only violating (P2) that can survive *any* cutting sequence.

The Buchholz Hydra generalizes the Hydra battle by allowing growth in height (violating (P3), although not arbitrarily, but basically restricted to an elementary Hydra battle. One could think of further generalizations along this line. What about a Hydra whose height growth is bounded by a Buchholz Hydra? We would conjecture that this Hydra is still doomed. But lacking a simple combinatorial proof for the Buchholz Hydra, we cannot find a proof for this conjecture.

Acknowledgements The particular style to draw the Hydra trees is due to Luccio and Pagli [23]. The literature review was done at the excellent library of Kyoto University during the author's visit in January 2006. We would like to thank Georg Moser for sending us an early draft of his recent paper on the Hydra rewriting battle.

References

1. Avigad, J.: Ordinal analysis without proofs. In: Sieg, W., Sommer, R., Talcott, C. (eds.) Reflections on the Foundations of Mathematics: Essays in Honor of Solomon Feferman. Lecture Notes in Logic, vol. 15, pp. 1–36. A.K. Peters Ltd., Wellesley (2002)
2. Beklemishev, L.D.: The worm principle. Technical Report 219, Department of Philosophy, University of Utrecht, Logic Group Preprint Series (2003)
3. Buchholz, W.: An independence result for $(\Pi^1_1 - CA) + BI$. Ann. Pure Appl. Log. **33**, 131–155 (1987)
4. Carlucci, L.: A new proof-theoretic proof of the independence of the Kirby-Paris' Hydra theorem. Theor. Comput. Sci. **300**(1–3), 365–378 (2003)
5. Cichon, E.A.: A short proof of two recently discovered independence results using recursion theoretic methods. Proc. Am. Math. Soc. **87**(4), 704–706 (1983)
6. Clote, P., McAloon, K.: Two further combinatorial theorems equivalent to the 1-consistency of Peano arithmetic. J. Symb. Log. **48**(4), 1090–1104 (1983)
7. Daly, B.: (no title) 17 August 2000. Also archived in the Mathematical Atlas at http://www.math.niu.edu/~rusin/known-math/00_incoming/goodstein
8. Dershowitz, N.: Trees, ordinals and termination. In: Proceedings of the 4th International Joint Conference CAAP/FASE on Theory and Practice of Software Development (TAPSOFT'93). Lecture Notes in Computer Science, vol. 668, pp. 243–250. Springer, New York (1993)
9. Dershowitz, N., Moser, G.: The Hydra battle revisited. In: Rewriting, Computation and Proof; Essays Dedicated to Jean-Pierre Jouannaud on the Occasion of his 60th Birthday. Lecture Notes in Computer Science, vol. 4600, pp. 1–27. Springer, New York (2007)

10. Devlin, K.: Devlin's angle: Infinity and intuition. In MAA Online: The Mathematical Association of America. January 2006. http://www.maa.org/devlin/devlin_01_06.html
11. Feferman, S., Friedman, H.M., Maddy, P., Steel, J.R.: Does mathematics need new axioms? Bull. Symb. Log. **6**(4), 401–446 (2006)
12. Gardner, M.: Mathematical games: Tasks you cannot help finishing no matter how hard you try to block finishing them. Sci. Am. 8–13 (1983)
13. Gentzen, G.: Die Widerspruchsfreiheit der reinen Zahlentheorie. Math. Ann. **112**, 493–565 (1936). Appendix: Galley proof of sections IV and V, Mathematische Annalen (1935) received on 11th August 1935. Translated as "The consistency of elementary number theory" in [27], pp. 132–213
14. Goodstein, R.J.: On the restricted ordinal theorem. J. Symb. Log. **9**, 33–41 (1944)
15. Hamano, M., Okada, M.: A relationship among Gentzen's proof-reduction, Kirby-Paris' Hydra game and Buchholz's Hydra game. Math. Log. Q. **43**, 103–120 (1997)
16. Hamano, M., Okada, M.: A direct independence proof of Buchholz's Hydra game on finite labeled trees. Arch. Math. Log. **37**, 67–89 (1998)
17. Ishihara, H.: Weak König's Lemma implies Brouwer's Fan Theorem. Notre Dame J. Form. Log. **47**(2), 249–252 (2006)
18. Jouannaud, J.-P.: Prroof and computation. In: Schwichtenberg, H. (ed.) NATO Series F: Computer and Systems Sciences, vol. 139, pp. 173–218. Springer, Heidelberg (1995)
19. Keisler, H.J.: Nonstandard arithmetic and reverse mathematics. Bull. Symb. Log. **12**(1), 100–125 (2006)
20. Kirby, L., Paris, J.: Accessible independence results for Peano arithmetic. Bull. Lond. Math. Soc. **14**, 285–293 (1982)
21. Lagarias, J.: The $3x + 1$ problem and its generalizations. Am. Math. Mon. **92**, 3–23 (1985)
22. Leadbetter, R.: Hydra. In: Encyclopedia Mythica (1999). http://www.pantheon.org/articles/h/hydra.html
23. Luccio, F., Pagli, L.: Death of a monster. ACM SIGACT News **31**(4), 130–133 (2000)
24. Paris, J.: Some independence results for Peano arithmetic. J. Symb. Log. **43**, 725–731 (1978)
25. The Lernean Hydra. In: Crane, G. (ed.) The Perseus Project. Tufts University, Department of the Classics (2000). http://www.perseus.tufts.edu/Herakles/hydra.html
26. Smullyan, R.M.: Trees and ball games. Ann. NY Acad. Sci. **321**, 86–90 (1979)
27. Szabo, M.E. (ed.): The Collected Papers of Gerhard Gentzen. North-Holland, Amsterdam (1969)
28. Tait, W.W.: Gödels reformulation of Gentzen's first consistency proof for arithmetic: The no-counterexample interpretation. Bull. Symb. Log. **11**(2), 225–238 (2005)
29. Tucholsky, K.: Ratschläge für einen schlechten Redner (Advice for a bad speaker). In: Zwischen gestern und morgen, pp. 95–96. Rowohlt Verlag, Hamburg (1952). English translation at http://www.nobel133.physto.se/Programme/tucholsky.htm
30. Wainer, S.S.: Accessible recursive functions. Bull. Symb. Log. **5**(3), 367–388 (1999)
31. Weiermann, A.: Classifying the phase transition of Hydra games and Goodstein sequences, 2006. Manuscript, available at http://www.math.uu.nl/people/weierman/goodstein.pdf